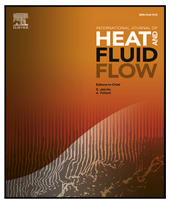
International Journal of Heat and Fluid Flow 117 (2026) 109988



Contents lists available at ScienceDirect

International Journal of Heat and Fluid Flow

journal homepage: www.elsevier.com/locate/ijhff



De-noising and super-resolution of fluid-flow velocity measurements by optimising a discrete loss (ODIL)

Stephen Terrington *, Mark Thompson, Kerry Hourigan

FLAIR, Department of Mechanical and Aerospace Engineering, Monash University, Clayton, 3800, VIC, Australia

ARTICLE INFO

Keywords:
Machine Learning
Fluid Mechanics
Physics-informed neural networks
optimising a discrete loss

ABSTRACT

This article presents a comparison between two techniques – optimising a discrete loss (ODIL) and physics informed neural networks (PINN) – for reconstructing the velocity field from low resolution and noisy planar PIV measurements. Both techniques are capable of accurately reconstructing the velocity and pressure fields from low resolution and noisy velocity measurements sampled from 2D numerical simulations. Both techniques provide reasonable reconstruction of the in-plane velocity components when provided with two-component velocity measurements in a single plane sampled from a 3D numerical simulation. However, ODIL generally over-fits to any noise in the measurement data, and therefore PINN achieves higher accuracy. While PINN can achieve a reconstruction more accurate than that of ODIL, PINN converges much slower than ODIL, requiring substantially more training epochs and walltime to produce results of similar accuracy to ODIL. Both methods are superior to statistical noise reduction approaches, such as low-pass filtering.

1. Introduction

This article compares the performance of two physics-informed machine-learning techniques – optimising a discrete loss (ODIL) and physics-informed neural networks (PINN) – for super-resolution and de-noising of fluid-flow velocity measurements. In the context of fluid mechanics, super-resolution (Fukami et al., 2023) refers to the reconstruction of a high-resolution velocity field from low-resolution velocity measurements. De-noising (Vétel et al., 2011) refers to the reconstruction of an accurate (i.e. low-noise) velocity field from noisy measurement data. Hence, the aim of the present work is to reconstruct a high-resolution and low-noise velocity field from low-resolution and noisy measurement data.

Data assimilation, such as adjoint-based data assimilation (ABDA) (He et al., 2018; Gillissen et al., 2019; He et al., 2022), is one approach to achieve de-noising and super-resolution. Here, the reconstructed velocity and pressure fields are constrained to exactly obey the governing physics. Then, a cost function representing the error between the reconstructed fields and the true solution is minimised using adjoint optimisation. This produces a high-resolution reconstruction which matches the measurement data as closely as possible while satisfying the known physics.

Alternatively, data-driven machine learning techniques have been employed to achieve super-resolution and de-noising (Fukami et al., 2023; Sofos et al., 2024; Sofos and Drikakis, 2025). In these approaches, a machine-learning model — such as an artificial neural network (ANN), convolution neural network (CNN), U-Net architecture, generative adversarial network (GAN), transformer or diffusion model (DM) – is trained on a large quantity of measurement data to learn the relationship between a low-resolution measurement field and the high-resolution reconstruction (Sofos and Drikakis, 2025). Once trained, the model can quickly compute the high-resolution velocity field when supplied with a low-resolution and noisy measurement.

The main downside of data-driven machine learning techniques is that a large quantity of training data is required. Moreover, known physics information is often not used in the reconstruction. Physics-informed neural networks (PINNs) were proposed by Raissi et al. (2019) to overcome these limitations in cases featuring limited data but which obey known physical laws. Here, the velocity and pressure fields are represented as a function of space and time using a neural network. The neural network is trained to minimise a loss function which includes both the error compared to the measurement data, and an error based on the known physical laws. This is conceptually similar to the adjoint-based data assimilation (He et al., 2018). However, for PINNs, the reconstructed solution only approximately satisfies the known physics, while in ABDA the reconstructed field exactly obeys known physics.

PINNs have previously been shown to be effective for the reconstruction of flow fields from sparse and noisy experimental velocity measurements (Wang et al., 2022; Hasanuzzaman et al., 2023; Xu et al.,

* Corresponding author.
E-mail address: stephen.terrington1@monash.edu (S. Terrington).

https://doi.org/10.1016/j.ijheatfluidflow.2025.109988
Received 15 April 2025; Received in revised form 4 July 2025; Accepted 14 July 2025
Available online 29 August 2025
0142-727X/© 2025 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

2022; Fan et al., 2023; Steinfurth and Wiess, 2024). However, the PINN technique was found to require a large number of iterations when the size of the neural network was increased (Wang et al., 2022), limiting its use in practical applications.

Optimising a discrete loss (ODIL) (Karnakov et al., 2023) has recently been proposed as an alternative to physics-informed neural networks (PINN) (Raissi et al., 2019) for physics-informed machine learning and data-assimilation tasks. In ODIL, the reconstructed velocity and pressure fields are represented using a discrete grid, rather than a neural network, and the finite-difference technique is used to compute the physics-based loss function. The total loss function (including both physics and data errors) is then optimised using the same gradient-based optimisers used for PINNs. The ODIL technique was found to converge much faster than the PINN technique (Karnakov et al., 2023), and is therefore a promising candidate for the reconstruction of fluid flows from experimental velocity measurements.

The present study compares the performance of these two techniques for the reconstruction of velocity and pressure fields from noisy and low resolution particle image velocimetry (PIV) measurements. The structure of this paper is as follows. First, in Section 2, we briefly summarise the PINN and ODIL techniques and describe the implementation used in this study. Then, in Section 3, we test these techniques on low resolution and noisy data sampled from a two-dimensional high-resolution numerical simulation. Next, in Section 4, the techniques are tested on two-dimensional two-component (2D2C) velocity fields, sampled from a high-resolution three-dimensional numerical simulation, imitating planar PIV measurements. Finally, in Section 5, the techniques are applied to experimental PIV measurements.

2. Methodology

A fluid flow has velocity and pressure fields $\boldsymbol{u}(\boldsymbol{x},t)$ and $p(\boldsymbol{x},t)$, respectively, where $\boldsymbol{x}$ is position, and t is time. We assume that low-resolution and noisy velocity measurements $\boldsymbol{v}_j$ are obtained,

$$v_{i,j} = u_i(\boldsymbol{x}_j, t_j) + \epsilon_{i,j}. \tag{1}$$

Here, subscript i represents the velocity component (either x, y or z), the subscript j refers to an individual measurement taken at position $\boldsymbol{x}_j$ and time t_j, and ϵ_i is the measurement error. The aim of the present study is to obtain a reconstruction of the true velocity field, $\boldsymbol{u}$, given only the measurements $\boldsymbol{v}_j$.

Assuming incompressible flow of a Newtonian fluid, the velocity and pressure are governed by the momentum and continuity equations,

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u}\cdot\nabla\boldsymbol{u} = -\nabla p + \frac{1}{Re}\nabla^2\boldsymbol{u}, \tag{2}$$

$$\nabla\cdot\boldsymbol{u} = 0. \tag{3}$$

Here, variables have been non-dimensionalised by a characteristic length L, velocity U, and the density ρ and dynamic viscosity μ of the fluid, and p is the kinematic pressure. Finally, $Re = \rho U L/\mu$ is the Reynolds number.

Both PINN and ODIL are physics-informed machine learning/data assimilation techniques. The known physics (Eqs. (2) and (3)) are combined with the measurement data (Eq. (1)) to obtain the reconstructed velocity and pressure fields.

A summary of the PINN technique is provided in Fig. 1(*a*). The reconstructed velocity and pressure fields are modelled as a neural network. The inputs to the neural network are the position $\boldsymbol{x}$ and time t, and the outputs are the velocity and pressure. In the present work, feed-forward multilayer neural networks are used. The neural network contains a number (N_{hl}) of hidden layers, with each containing N_{nrn} neurons per hidden layer. Each neuron computes a linear weighted average of the outputs of the previous layer, which is then passed through a non-linear activation function. For the present study, the hyperbolic tangent function is used for the activation function, since PINN requires a continuously differentiable activation function. Such a neural network can represent any arbitrary function. In this case, the velocity and pressure as a function of position and time are represented using a neural network.

The free parameters of the neural network – namely the weights and biases of each neuron – are determined by optimising a loss function using a gradient descent algorithm. As is common in PINN studies (e.g. Wang et al., 2022), we use a combination of the ADAM and L-BFGS optimisers in the present work. The ADAM optimiser is a stochastic gradient-based optimiser, which quickly converges to an approximate solution, but is slow to converge to an accurate solution. The L-BFGS optimiser converges towards a more accurate solution than ADAM, but is not guaranteed to converge to a global optimum. Therefore, ADAM is used for the first N_{adam} epochs, before switching to the L-BFGS optimiser (Wang et al., 2022). For all examples in this study, the weights of the neural network were initialised using the LeCun normal initialiser, while the biases were set to zero.

The loss function includes both the measurement loss,

$$L_{meas} = \left|\left|\boldsymbol{v}_i - \boldsymbol{u}_i\right|\right|_2, \tag{4}$$

and the physics losses

$$L_{mom} = \left|\left|\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u}\cdot\nabla\boldsymbol{u} + \nabla p - \frac{1}{Re}\nabla^2\boldsymbol{u}\right|\right|_2, \tag{5}$$

$$L_{cont} = \left|\left|\nabla\cdot\boldsymbol{u}\right|\right|_2, \tag{6}$$

where $\|\ \|_2$ is the $L2$ norm, or root-mean-square (r.m.s.) error. The physics losses are evaluated at a set of N_c interior collocation points, randomly sampled using Latin hypercube sampling. Finally, the total loss function is a weighted sum of the physics and measurement losses:

$$L_{total} = W_{meas}L_{meas} + W_{phys}(L_{mom} + L_{cont}), \tag{7}$$

where W_{meas} and W_{phys} are weights for the measurement and physics loss, respectively. Minimising the total loss function obtains a reconstructed velocity field that approximately matches the measurement data, while also approximately satisfying the momentum and continuity equations. For further details on the PINN method, refer to Raissi et al. (2019).

The hyperparameters for PINN include the loss function weights (W_{meas} and W_{phys}), the number of hidden layers N_{hl}, the number of neurons per hidden layer N_{nrn}, the number of interior collocation points N_c, the learning rate for ADAM l_r, the number of epochs for which ADAM optimisation is used N_{adam}, and the L-BFGS parameter m_{lbfgs}, which represents the number of position and gradient updates stored in memory when using the L-BFGS optimiser. The hyperparameters used for each test case in the present study are presented at the start of each section of this manuscript.

A summary of the ODIL technique is presented in Fig. 1(*b*). For ODIL, the reconstructed velocity and pressure field are represented using a discrete grid, as used in existing computational fluid dynamics schemes. We assume that the measurement data are obtained on a regular Cartesian grid, as is the case for PIV measurements. Then, the reconstructed fields are represented on a fine Cartesian grid. The loss functions are defined as in Eqs. (4)–(6). However, derivatives for the physics losses are evaluated using second order finite difference schemes. A second order upwind scheme is used for the advection terms, while a central differencing scheme is used for the remaining terms.

Additionally, regularisation terms are added to the loss function to promote smooth solutions (Karnakov et al., 2023). Unless otherwise specified, a fourth-order smoothing was used, since the momentum equation includes terms up to second order:

$$L_{reg} = \left|\left|\frac{\partial^4 u}{\partial x^4} + \frac{\partial^4 u}{\partial y^4}\right|\right|_2 + \left|\left|\frac{\partial^4 v}{\partial x^4} + \frac{\partial^4 v}{\partial y^4}\right|\right|_2 + \left|\left|\frac{\partial^4 p}{\partial x^4} + \frac{\partial^4 p}{\partial y^4}\right|\right|_2. \tag{8}$$

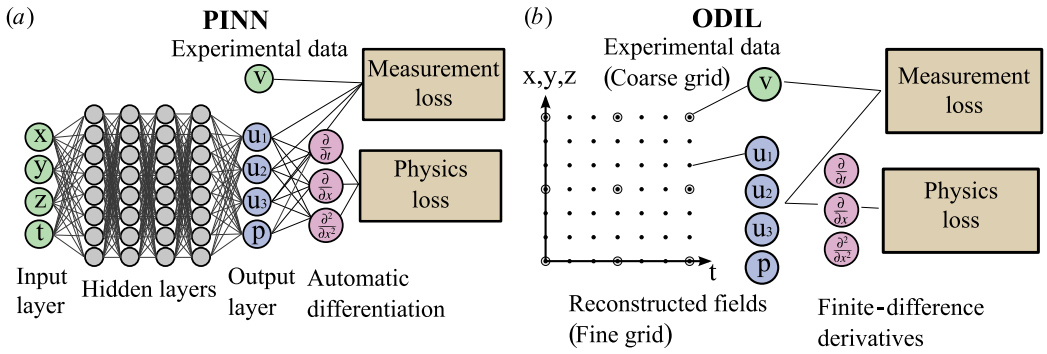


Fig. 1. Diagrams summarising the (*a*) PINN and (*b*) ODIL techniques.

The fourth-order derivatives were computed using a second-order accurate central finite difference approximation, using a 5-point stencil in each direction.

The free parameters are the discrete velocity and pressure values at each grid-point. Unless otherwise stated, these are initialised by interpolating from the low-resolution measurement data as an initial condition. Then, the free parameters are optimised by minimising the total loss function,

$$L_{total} = W_{meas} L_{meas} + W_{phys}(L_{mom} + L_{cont}) + W_{reg} L_{reg}. \tag{9}$$

These are optimised by minimising the total loss function using the L-BFGS optimiser. For further details on the ODIL technique, see Karnakov et al. (2023).

The hyperparameters for ODIL include the loss function weights (W_{meas}, W_{phys} and W_{reg}), the spatial and temporal resolution of the high-resolution reconstruction (Δx_r, Δy_r and Δt_r), and the L-BFGS parameter m_{lbfgs}. The hyperparameters used for each test case in the present study are presented at the start of each section in this manuscript.

For the present work, both ODIL and PINN were implemented using the Python software ODIL, which is available at https://github.com/cselab/odil (Karnakov et al., 2023). Python scripts for running the various example studies are available on GitHub (https://github.com/Sterring/pinn-odil-compare). All examples were performed on a desktop workstation equipped with an NVIDIA GeForce RTX 4090 graphics card (24 GB), and an Intel® Core™ i7-5930K CPU (3.50 GHz, 6 cores).

In summary, both ODIL and PINN obtain reconstructed velocity and pressure fields by minimising a loss function that includes both measurement loss and physics loss. For PINN, the reconstructed velocity is represented as a neural network, and derivatives are computed using automatic differentiation. For ODIL, the reconstructed velocity is represented using a discrete grid and finite difference techniques are used to represent derivatives.

3. 2D numerical simulation data

We first compare the performances of PINN and ODIL in reconstructing low resolution and noisy data sampled from high resolution 2D numerical simulation data. We use the high-resolution cylinder wake presented in Raissi et al. (2019), which is publicly available on gitHub (https://github.com/maziarraissi/PINNs). These data were obtained using the spectral/hp-element solver NekTar, for the wake behind a circular cylinder at $Re = 100$. For further details, see Raissi et al. (2019).

The high-resolution reference solution is represented using a regular grid with resolution $\Delta x = 0.0707$, $\Delta y = 0.0816$ and $\Delta t = 0.1$. Synthetic measurement datasets were created by down-sampling the high resolution numerical data, and adding Gaussian noise with standard deviation σ_ϵ. In addition to lowering the resolution, the domain of the measurement data was restricted to $x_{min} < x < x_{max}$, $y_{min} < y < y_{max}$ and $t_{min} < t < t_{max}$. Unless otherwise specified, these parameters were set to $x_{min} = 1$, $x_{max} = 8$, $y_{min} = -2$, $y_{max} = 2$, $t_{min} = 0$ and $t_{max} = 5$. Note that this domain does not include the cylinder, which is centred at $x = 0$. To accurately capture the boundary layer of the cylinder using ODIL would require a much finer grid than is practical using a regular Cartesian grid. Combining ODIL with a different discretisation scheme, such as the finite volume or finite element method, may allow the boundary layers to be captured, which would be necessary to compute the forces applied to the cylinder. However, this is beyond the scope of the present study.

Table 1
List of hyperparameters used for both ODIL and PINN to reconstruct synthetic measurements of the 2D cylinder wake.

ODIL		PINN	
W_{meas}	1	W_{meas}	1
W_{phys}	5	W_{phys}	1
W_{reg}	10^{-4}	N_{hl}	6
Δx_r	0.0707	N_{nrn}	20
Δy_r	0.0816	N_c	10,000
Δt_r	0.05	l_r	10^{-2}
m_{lbfgs}	20	N_{adam}	500
		m_{lbfgs}	20

Fig. 2(*a*) presents an example synthetic dataset with $\sigma_\epsilon = 0.1$ and $\Delta x = 0.283, \Delta y = 0.327$ and $\Delta t = 0.4$. The pressure field from the numerical simulation was not included in the synthetic measurement datasets. Subfigures (*b*) and (*c*) present a comparison between the high-resolution reference solution, and the reconstructions obtained using both PINN and ODIL. The hyperparameters for ODIL and PINN are listed in Table 1. Further details on the selection of hyperparameters is provided in Sections 3.1 & 3.2. Both techniques are able to satisfactorily recover the original velocity fields from noisy measurement data (subfigure *b*). Moreover, both techniques also accurately reconstruct the pressure field (subfigure *c*), even though measurement data for the pressure were not supplied. To enable a fair comparison, the pressure fields were first normalised to have zero mean by subtracting a constant pressure. This is necessary as pressure is only uniquely defined up to an arbitrary constant.

For this figure, as well as the remainder of the paper, vorticity was computed using second-order accurate central finite-differencing from the discrete velocity fields (either the low-resolution synthetic measurement data; the high-resolution reference solution; the ODIL reconstruction; or the PINN reconstruction). For PINN, vorticity was computed by post-processing the discrete velocity fields obtained by evaluating the neural network on a discrete grid with equivalent resolution to the reference solution.

Fig. 3 presents a comparison of the convergence between ODIL and PINN techniques. This figure presents the L2 norm of error between the reconstructed velocity and pressure fields, and the high-resolution

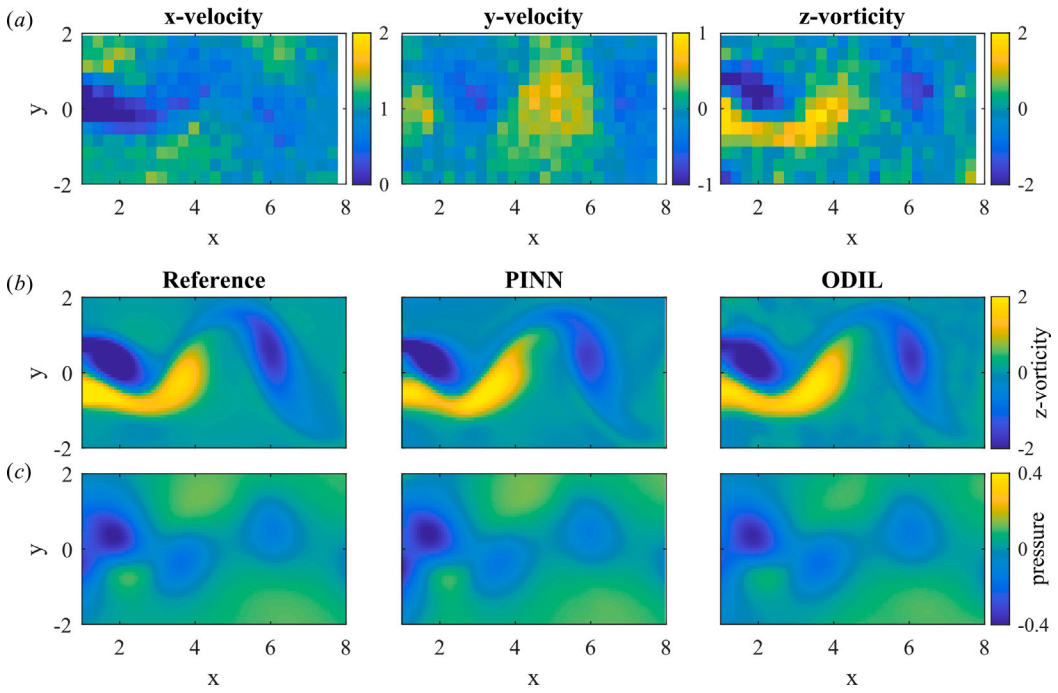


Fig. 2. (*a*) Synthetic low-resolution and noisy velocity measurements sampled from high resolution data, and (*b*, *c*) comparison between the reconstructed (*b*) vorticity and (*c*) pressure using both ODIL and PINN, to the high resolution reference solution. An animated version of this figure is available in the online supplementary material.

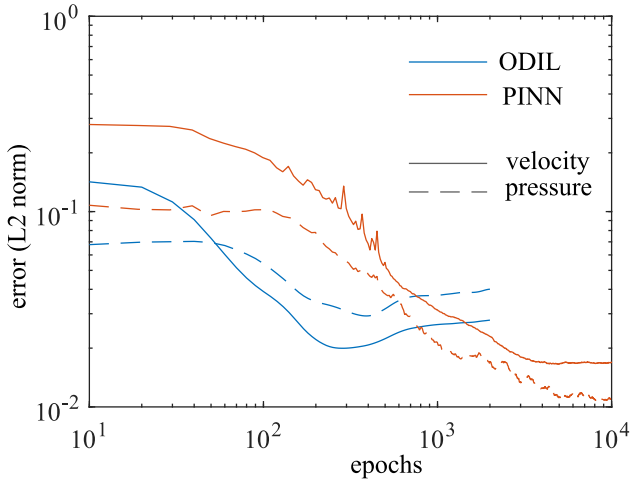


Fig. 3. Comparison between the convergence history of ODIL and PINN, showing the L2 norm of the error between the reconstructed velocity and pressure fields and the high-resolution reference solution.

reference solution. Specifically, the solution error for velocity and pressure are computed as:

$$E_u = \frac{1}{2}\left(\sqrt{\frac{1}{N_{hr}}\sum^{N_{hr}}(u_x - \hat{u}_x)^2} + \sqrt{\frac{1}{N_{hr}}\sum^{N_{hr}}(u_y - \hat{u}_y)^2}\right), \tag{10}$$

$$E_p = \sqrt{\frac{1}{N_{hr}}\sum^{N_{hr}}(p - \hat{p})^2}, \tag{11}$$

where u_x, u_y and p represent the reconstructed velocity and pressure fields, $\hat{u}_x$, $\hat{u}_y$ and $\hat{p}$ represent the reference solution and the summation is over each point in the high-resolution reference solution, with N_{hr} being the total number of points in the reference solution. The various parameters for both ODIL and PINN were selected to achieve a minimum solution error, as described in Sections 3.1 & 3.2. The minimum solution error achieved using ODIL was 2.49×10^{-2} (2.00×10^{-2} for velocity, 2.92×10^{-2} for pressure), requiring 370 epochs of training, and 66 s of walltime. After this, the solution error increases due to overfitting. PINN converges slower than ODIL, taking 1060 epochs (walltime = 155 s) to achieve a solution error equal to ODIL. However, PINN is less prone to overfitting, and achieves a lower overall solution error of 1.38×10^{-2} (1.67×10^{-2} for velocity and 1.08×10^{-2} for pressure) in 9318 epochs (walltime = 1099 s). Therefore, ODIL converges faster than PINN. However, ODIL is prone to overfitting, and PINN achieves a lower minimum solution error.

3.1. ODIL

We now discuss the effect of various parameters on the performance of ODIL. These parameters include: the physics, data and regularisation weights (W_{phys}, W_{dat} and W_{reg}); the resolution of the reconstructed fields (Δx_r, Δy_r, Δt_r); the loss function norm; and the initial conditions. The optimal parameters were found to be: $W_{phys} = 5$, $W_{dat} = 1$, $W_{reg} = 10^{-4}$, $\Delta x_r = 0.0707$, $\Delta y_r = 0.0816$ and $\Delta t_r = 0.05$. Moreover, the lowest solution error was achieved using the L2 norm for all terms in the loss function, and by using the measurement data as an initial condition. Unless otherwise stated, these parameters are used throughout this section.

Fig. 4(*a*) plots the solution error calculated according to Eqs. (10) and (11) against number of training epochs for ODIL with these parameters. A minimum solution error of 2.49×10^{-2} (2.00×10^{-2} for velocity, 2.92×10^{-2} for pressure) is achieved in 370 epochs of training, requiring 66 s of walltime. Further training results in an increase in error, indicating over-fitting. Subfigure (c) plots the vorticity and pressure solutions after 100, 400 and 10,000 epochs. After 100 epochs, the wake vortices are not fully resolved, resulting in low vorticity and pressure magnitudes. After 400 training epochs, the wake structures are accurately resolved. At 10,000 training epochs, overfitting manifests as added noise near the inlet.

To test for over-fitting in situations where the true solution is not known, we split the measurement data into training (90% of data) and validation (10% of data) sets. Only the training data are supplied to ODIL/PINN. Fig. 4(*b*) presents the convergence of the mean squared

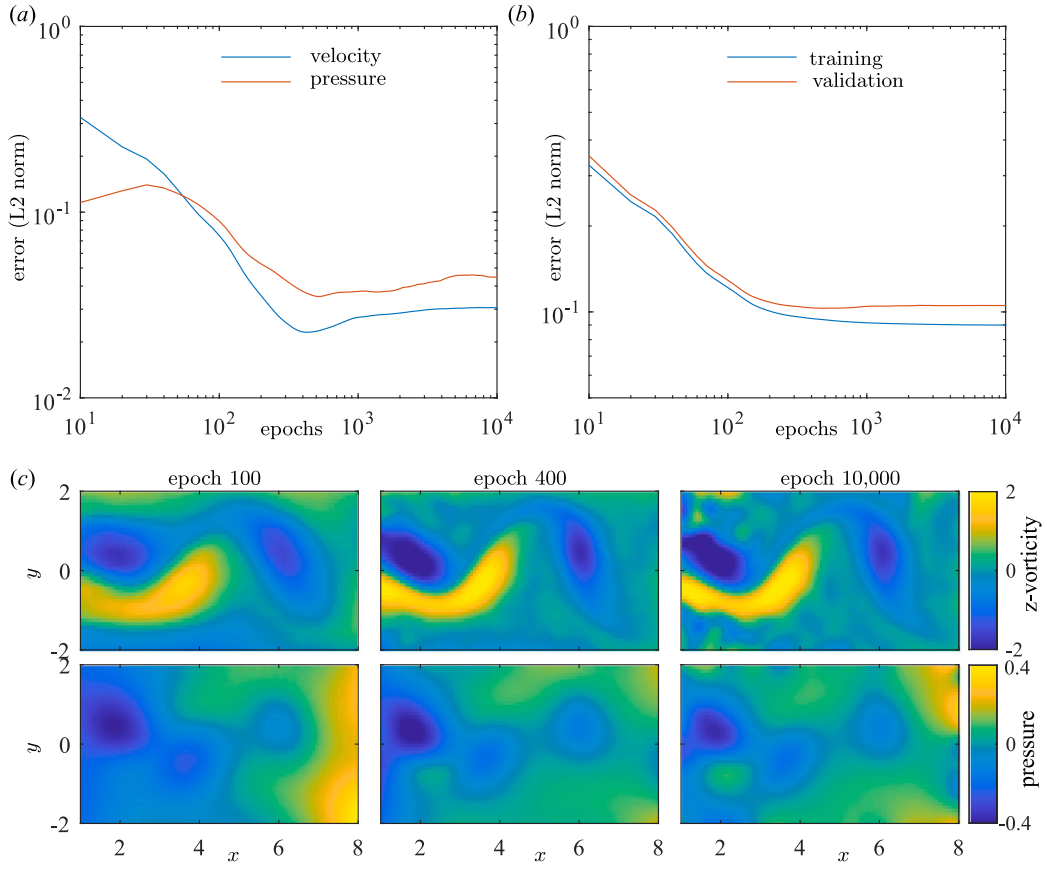


Fig. 4. Convergence history of ODIL, showing (*a*) the L2 norm of the error between the ODIL reconstruction and the reference solution and (*b*) the validation and training errors. (*c*) Reconstructed vorticity and pressure fields after 100, 400 and 10,000 epochs.

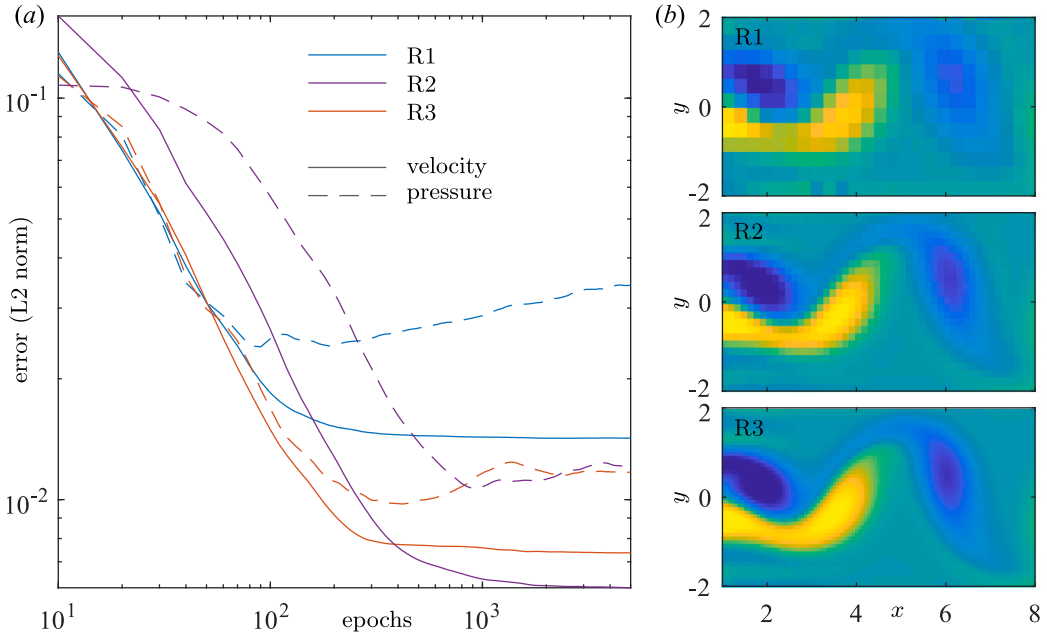


Fig. 5. Effect of reconstruction resolution on the solution error, showing (*a*) convergence history of the solution error, and (*b*) reconstructed vorticity fields at each resolution. The reconstruction resolutions are: $\Delta x_r = 0.2828$, $\Delta y_r = 0.3265$ and $\Delta t_r = 0.2$ for R1; $\Delta x_r = 0.1414$, $\Delta y_r = 0.1633$ and $\Delta t_r = 0.1$ for R2; $\Delta x_r = 0.0707$, $\Delta y_r = 0.0816$ and $\Delta t_r = 0.05$ for R3.

error between the reconstructed solution and both the training and validation dataset. The training and validation errors are computed as:

$$E_{train} = \frac{1}{2}\left(\sqrt{\frac{1}{N_{train}}\sum^{N_{train}}(u_x - v_x)^2} + \sqrt{\frac{1}{N_{train}}\sum^{N_{train}}(u_y - v_y)^2}\right), \quad (12)$$

$$E_{val} = \frac{1}{2}\left(\sqrt{\frac{1}{N_{val}}\sum^{N_{val}}(u_x - v_x)^2} + \sqrt{\frac{1}{N_{val}}\sum^{N_{val}}(u_y - v_y)^2}\right), \quad (13)$$

where u and v are the measurement data and reconstructed fields, respectively, while N_{train} and N_{val} are the number of points in the training and validation sets, respectively. While the training error decreases monotonically with number of training epochs, the validation error begins to increase after 490 epochs, which is slightly greater than the 370 epochs required to minimise the error with respect to the reference solution. Hence, the error compared to the validation data provides a reasonable test for overfitting, in the absence of a reference solution.

Fig. 5 shows the effect of reconstruction resolution on solution accuracy, for noiseless ($\sigma_\epsilon = 0$) data. The labels R1, R2 and R3 correspond to increasing reconstruction resolutions as follows. For R1: $\Delta x_r = 0.2828$, $\Delta y_r = 0.3265$ and $\Delta t_r = 0.2$; for R2: $\Delta x_r = 0.1414$, $\Delta y_r = 0.1633$ and $\Delta t_r = 0.1$; for R3: $\Delta x_r = 0.0707$, $\Delta y_r = 0.0816$ and $\Delta t_r = 0.05$. R3 has the same spatial resolution as the reference solution, while R2 and R1 are coarser by a factor of 2 and 4, respectively. The temporal resolution was selected to keep the Courant number below unity. Increasing the reconstruction resolution results in a higher accuracy of the converged solution, at the cost of increased number of training epochs and training time. The R2 resolution is sufficient to achieve below 1% solution accuracy. However, the R3 resolution is used in the remainder of this section, as this is the same resolution as the reference solution.

As discussed in Section 2, the L_2 norm (mean squared error) is used for all loss functions. The L_1 norm (mean absolute error) was also considered, as it typically less sensitive to outliers (e.g. measurement noise). Fig. 6 compares the solutions obtained using the L_2 norm for all terms (MSE), the L_1 norm for all terms (MAE), and the L_1 norm only for the data loss (MSE_P). Subfigure (a) plots the convergence history of the solution error, while subfigure (b) plots vorticity contours for the trained solutions (i.e. at a number of epochs corresponding to minimum solution error). Using MAE for all terms results in a poor solution error and slow convergence, showing artefacts in the vorticity plot. Using the L_2 norm for all terms (MSE) results in the lowest solution error, while using the L_1 norm for only the data loss (MSE_P) yields a slightly worse reconstruction. Therefore, the L_2 is used exclusively in the remainder of this study.

We now consider the effects of the physics and data weights. For convenience, the data weight $W_{dat} = 1$ is kept at unity, and only the physics weight W_{phys} is varied. Fig. 7(a) plots the convergence of the solution error for several different physics weights, while subfigure (b) plots vorticity contours obtained for different physics weights. Increasing W_{phys} results in a smoother solution with less noise, as the relative importance of the data loss is reduced. The minimum solution error is achieved at both $W_{phys} = 2$ and 5 (which differ by less than 0.2%). A value of $W_{phys} = 5$ was chosen, as it produces smoother vorticity contours.

Fig. 8(a) plots the convergence history for various regularisation weights, while subfigure (b) plots pressure contours obtained using each weight. Using a small regularisation weight $W_{reg} = 10^{-4}$ results in a slightly lower solution error, compared to $W_{reg} = 0$. As shown in subfigure (b), the reconstructed pressure fields are not smooth for $W_{reg} = 0$ and $W_{reg} = 10^{-5}$, leading to lower solution accuracy. Further increasing the weight to $W_{reg} = 10^{-3}$ results in a much slower training. Therefore, $W_{reg} = 10^{-4}$ was used in the remainder of this section.

One additional benefit of ODIL over PINN is that one can more easily set the initial solution prior to training. Fig. 9 plots the convergence history and vorticity plots for two cases. The case IC-0 uses all zeros for the initial velocity and pressure fields, while IC-M uses the noisy measurement data as an initial condition. Using the initial measurement data as an initial condition results in a lower minimum solution error, and a faster training time.

Table 2
Variation of the number of points in the training dataset and the walltime per epoch with the resolution of the measurement data and the extent of the temporal domain (t_{max}).

Δx	Δy	Δt	t_{max}	N_{dat}	Walltime/epoch
0.2828	0.3265	0.4	1	975	126 ms
0.2828	0.3265	0.4	5	4225	140 ms
0.2828	0.3265	0.4	10	8125	160 ms
0.2828	0.3265	0.4	20	16,250	207 ms
0.1414	0.1633	0.2	5	31,250	144 ms
0.5657	0.6531	0.8	5	637	155 ms

We now consider the effect of the amount of training data on the performance. The quantity of training data may be increased either by increasing the spatial resolution of the measurement data, or by increasing the size of the domain. We first consider the size of the time domain, by varying the parameter t_{min}. Fig. 10(a) presents the convergence histories of solution error (calculated according to Eqs. (10) & (11)). This comparison was performed with a fixed resolution of measurement data ($\Delta x = 0.283, \Delta y = 0.327$ and $\Delta t = 0.4$), so that the number of samples of training data increases with t_{min}, as shown in Table 2. Moreover, the measurement data had added noise with $\sigma_\epsilon = 0.1$. The minimum solution error achieved decreases as t_{min} is increased, indicating that a more accurate solution is achieved with a larger temporal domain for fixed resolution. Both the number of training epochs required to achieve a minimum solution error and the walltime required per epoch increase as t_{max} increases.

Fig. 10(b) plots the variation of solution error against time. Here, the errors are calculated as in Eqs. (10) & (11), but keeping the flow-time fixed. The solution errors are largest near the initial flow time, then decrease to a nearly constant level, with some evidence of a periodic variation in the error. There is also a small increase in error near the final time. The errors are likely largest near the initial condition and domain boundaries due to no initial or boundary conditions being applied (since these are not known from the measurement data set). The average solution error decreases with increased domain size, as the influence of these errors becomes less significant.

We now consider the effect of the resolution of the training data, while keeping the size of the domain fixed. As shown in Table 2, increasing the resolution of the training data does not substantially decrease the walltime required per epoch. This is since the number of data points required to represent the high-resolution solution is much larger than the number of points in the training data set, and is independent of the resolution of the measurement data.

Fig. 11 plots vorticity fields for reconstructed solutions for a range of data resolutions, and measurement noises. In general, ODIL displays a remarkable ability to recover the vortex structure from low resolution and noisy data. The accuracy of reconstruction decreases as the amount of noise is increased, or resolution of the measurement data is decreased. Even with substantial noise ($\sigma_\epsilon = 0.5$), for which the vortex structure is not clearly visible in the original measurement data, the vortex structure is clearly recovered in the reconstruction, albeit with significant noise. The minimum solution errors for each of the cases shown in Fig. 11 are provided in Table 3. As either the amount of noise is increased, or the resolution of the measurement data is increased, the accuracy of the ODIL reconstruction is reduced. In all cases, however, ODIL substantially reduces the measurement noise and increases the spatial resolution, compared to the initial measurement data.

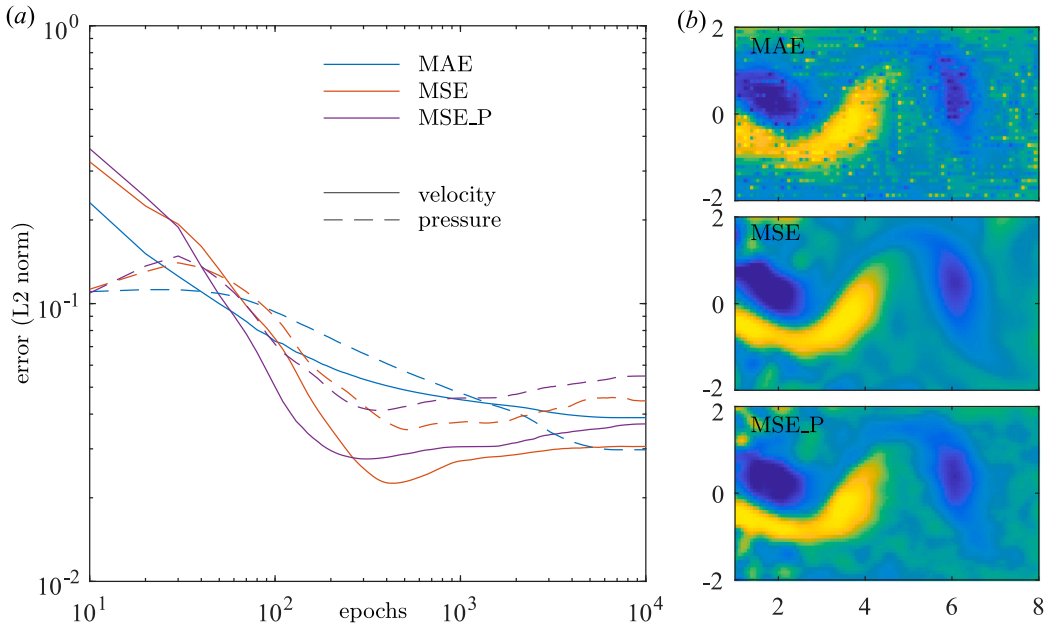


Fig. 6. Effect of loss function definition on (*a*) convergence history of the solution error and (*b*) reconstructed vorticity fields. The labels are as follows: MAE is the mean absolute error (L1 norm); MSE is the mean-squared error (L2 norm); MSE_P uses the L1 norm for the data loss, and the L2 norm for all other terms.

Table 3
L2 norm of the error between the ODIL and PINN reconstructions and the high-resolution reference solution, as well as the number of training epochs and training wall-time, for various resolution and noise of the measurement data.

Δx	Δy	Δt	σ_ϵ	E	E_u	E_p	Epochs	Walltime (s)
0.1414	0.1633	0.2	0.01	7.05×10^{-3}	6.39×10^{-3}	7.70×10^{-3}	520	78
			0.05	9.53×10^{-3}	7.87×10^{-3}	1.12×10^{-2}	740	105
			0.10	1.49×10^{-2}	1.19×10^{-2}	1.79×10^{-2}	1000	144
			0.50	6.95×10^{-2}	5.33×10^{-2}	8.58×10^{-2}	1000	141
0.2828	0.3265	0.4	0.01	9.38×10^{-3}	8.65×10^{-3}	1.01×10^{-2}	670	99
			0.05	1.48×10^{-2}	1.37×10^{-2}	1.60×10^{-2}	560	84
			0.10	2.79×10^{-2}	2.15×10^{-2}	3.43×10^{-2}	280	46
			0.50	1.59×10^{-1}	8.14×10^{-2}	2.37×10^{-1}	250	42
0.5657	0.6531	0.8	0.01	1.91×10^{-2}	1.91×10^{-2}	1.91×10^{-2}	1000	156
			0.05	2.98×10^{-2}	3.12×10^{-2}	2.83×10^{-2}	650	104
			0.10	4.21×10^{-2}	4.55×10^{-2}	3.87×10^{-2}	370	62
			0.50	1.11×10^{-1}	1.19×10^{-1}	1.03×10^{-1}	170	33

3.2. PINN

We now discuss the effects of the neural network architecture on the performance of PINN. We vary three parameters: the number of neurons per hidden layer (N_{nrn}), the number of hidden layers (N_{hl}) and the number of interior collocation points (N_c).

Fig. 12(*a*) shows the effect of varying the number of neurons per hidden layer on the convergence history of PINN for a case with data resolution $\Delta x = 0.283, \Delta y = 0.327$ and $\Delta t = 0.4$ and with either no added noise ($\sigma_\epsilon = 0$), or with added noise ($\sigma_\epsilon = 0.1$). This figure shows that $N_{nrn} = 10$ is not able to accurately reproduce the reference solution, while $N_{nrn} \geq 20$ is sufficient to reproduce the original solution. $N_{nrn} = 20$ converges slightly faster than $N_{nrn} = 40$ and 80, and therefore $N_{nrn} = 20$ is used for the remainder of this section.

Fig. 12(*b*) shows the effect of varying the number of hidden layers (N_{hl}). This figure demonstrates that $N_{hl} = 6$ results in the lowest solution error, and the fewest training epochs to reach the converged solution. Therefore, a network containing 6 hidden layers, and 20 neurons per hidden layer is optimum for this problem.

Fig. 13 shows the effect of varying the number of interior collocation points on the convergence history of PINN. The case $N_c = 0$ corresponds to a pure PINN interpolation with no physics information. For pure interpolation, a minimum error of 2.67×10^{-2} for the velocity is achieved, while pressure is not obtained. Increasing the number of interior collocation points results in improved accuracy of the reconstructed velocity fields, and improved reconstruction of the pressure fields. While $N_c = 1000$ achieves satisfactory results for both velocity and pressure, $N_c = 10{,}000$ results in a lower overall solution error, due to improved reconstruction of the pressure field, and therefore $N_c = 10{,}000$ is used for the remainder of this study.

Fig. 14 plots the variation of solution error against (*a*) flow time t and (*b*) streamwise position x, for both ODIL and PINN. For both ODIL and PINN, the solution errors are largest near the initial condition ($t = 0$) and near the domain boundaries ($x = 1$ and $x = 8$).

We now compare the performance of different optimisers (either ADAM or L-BFGS) on the convergence of both PINN and ODIL. For this test, the data resolution was set to $\Delta x = 0.2828$, $\Delta y = 0.3265$ and $\Delta t = 0.4$, with either noiseless ($\sigma_\epsilon = 0$) or noisy ($\sigma_\epsilon = 0.1$) data. Fig. 15 presents the convergence history of solution error against number of training epochs for both PINN and ODIL, with either noiseless (*a*) or noisy (*b*) measurement data. Optimisations were performed using

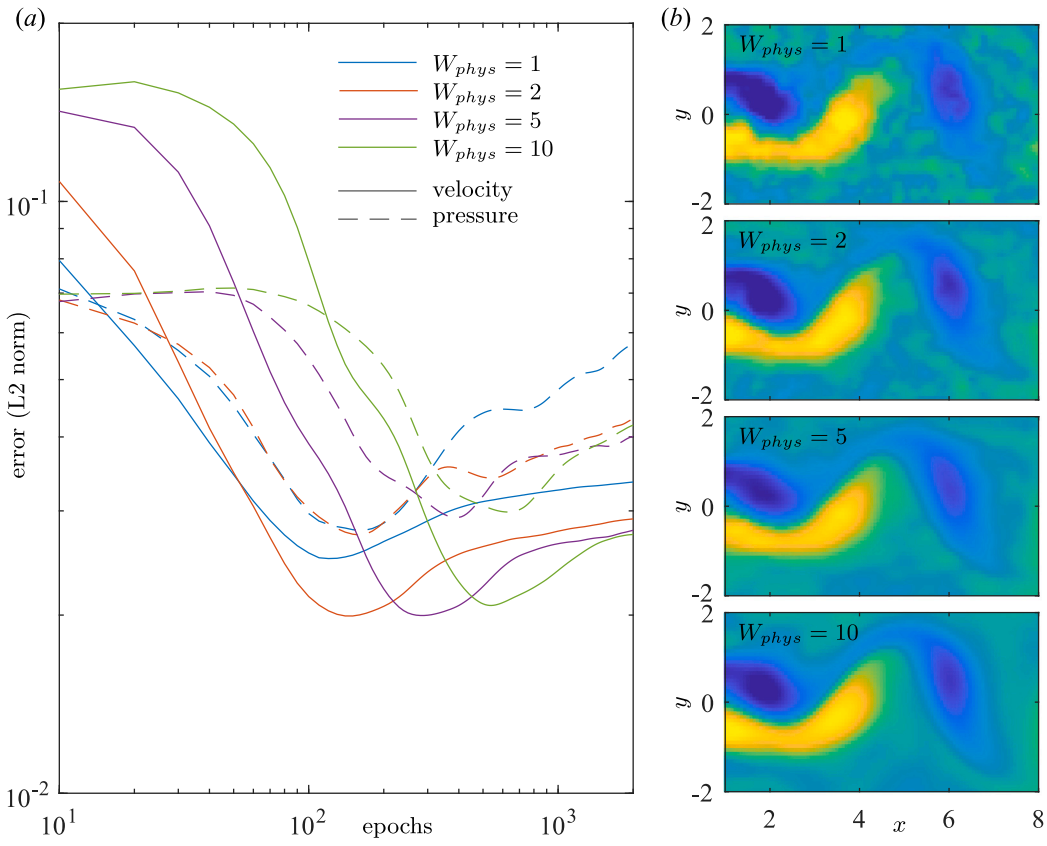


Fig. 7. Effect of physics weight on (*a*) convergence of the solution error, and (*b*) reconstructed vorticity fields.

either the ADAM or the L-BFGS optimisers. A third case used the ADAM optimiser for the initial stages of optimisation (100 epochs for ODIL, 500 epochs for PINN) before switching to the L-BFGS optimiser, which is labelled ADAM+L-BFGS in Fig. 15. The learning rate for ADAM was set to 1×10^{-2}, as this achieved fastest convergence. When using a reduced learning rate of 1×10^{-3}, ADAM was significantly slower than L-BFGS for both ODIL and PINN.

For ODIL, the ADAM optimiser with a learning rate $l_r = 0.01$ initially converges at a similar rate to the L-BFGS optimiser. However, a lower solution error can be achieved using the L-BFGS optimiser. Moreover, using the ADAM optimiser for the initial stages of the optimisation is slightly slower and less accurate than using the L-BFGS optimiser for the entire optimisation. For noiseless data, ODIL using L-BFGS requires 660 epochs (walltime = 98 s) to reach a solution error below 10^{-2}, while ADAM+L-BFGS requires 790 total epochs (walltime = 122 s). For data with added noise $\sigma_\epsilon = 0.1$, ODIL with L-BFGS reaches a minimum solution error of 2.16×10^{-2} using 71 s of walltime (460 epochs), while ADAM+L-BFGS has a slightly higher minimum error of 2.27×10^{-2}, in a similar training time (75 s; 450 epochs). Therefore, for this case L-BFGS is the preferred optimiser for ODIL.

For PINN, however, using pure L-BFGS has the slowest convergence, requiring 640 s of walltime (5319 epochs) to reduce the solution error below 0.01. The ADAM optimiser with a learning rate of 0.01 required only 571 s (7869 epochs) to reach the same solution error, while ADAM+L-BFGS required 501 s (4128 epochs). For data with added noise $\sigma_\epsilon = 0.1$, the L-BFGS optimiser required 280 s of walltime (2249 epochs) to reduce the solution error below 2.16×10^{-2} (the minimum error achieved using ODIL). The ADAM optimiser required 181 s (2180 epochs) to achieve the same error, while ADAM+L-BFGS required 249 s (2128 epochs). Using ADAM+L-BFGS, a minimum error of 1.65×10^{-2} is obtained in 582 s of walltime (4958 epochs), while using ADAM alone required only 479 s of walltime (6619 epochs) to achieve the same error. Therefore, ADAM uses less walltime than ADAM+L-BFGS for noisy data, but is slower for low-noise data. Since ADAM+L-BFGS can produce a lower solution error for low-noise data, and noting that the differences in walltime required are relatively modest, this combination of optimisers is used for PINNs in the remainder of this study.

Finally, Fig. 16 presents the vorticity fields reconstructed using PINN for a range of measurement data resolutions, and with $\sigma_\epsilon = 0.1$. The reference solution is accurately reconstructed from noisy and low resolution data. Table 4 lists the solution errors obtained in each case. Compared with the ODIL cases (Table 3), PINN achieves a more accurate reconstruction for the same quality of measurement data. Table 4 also lists the number of epochs and walltime required to reduce the error to below the corresponding ODIL case. In all cases, PINN requires much more training time to achieve a comparable error to ODIL. Figs. 12 and 16 show that reducing the complexity of the neural network will not achieve comparable error to ODIL in similar training time.

4. 3D numerical simulation data

In many cases, such as planar PIV experiments, only two velocity components are obtained in a two-dimensional measurement plane, sampled from a fully three-dimensional flow (two-dimensional two-component or 2D2C measurement). In this section, we compare the

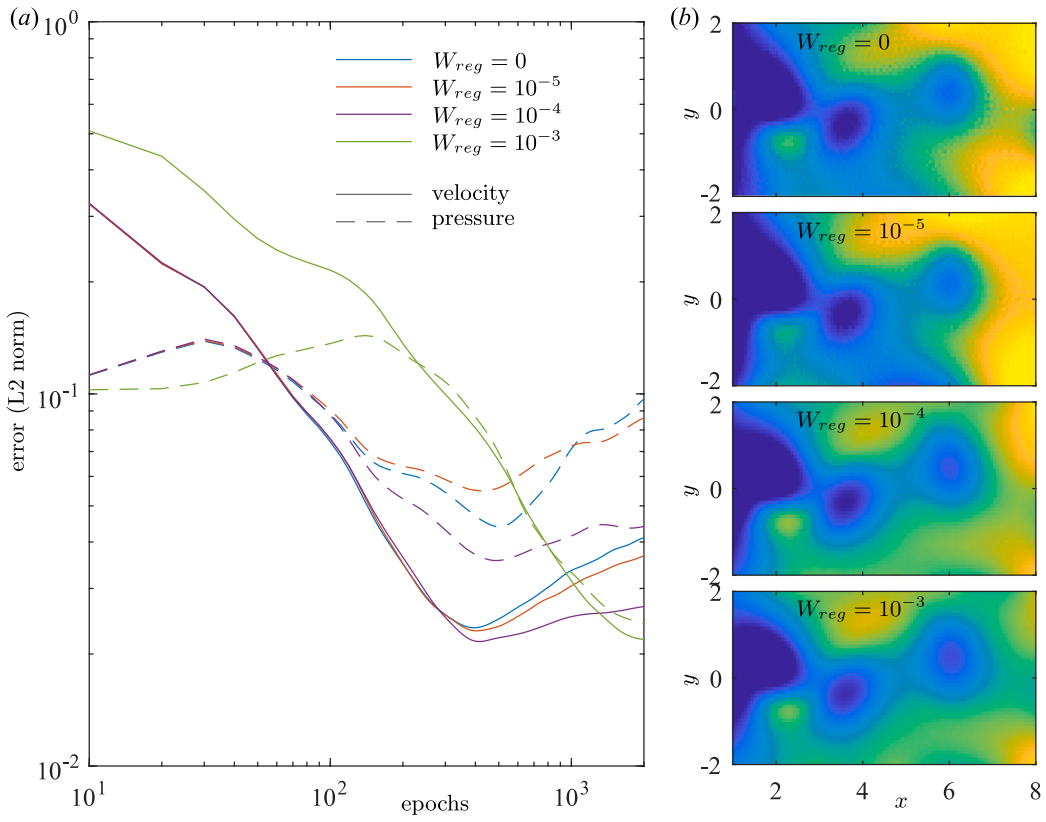


Fig. 8. Effect of regularisation weight on (*a*) convergence of the solution error, and (*b*) reconstructed pressure fields.

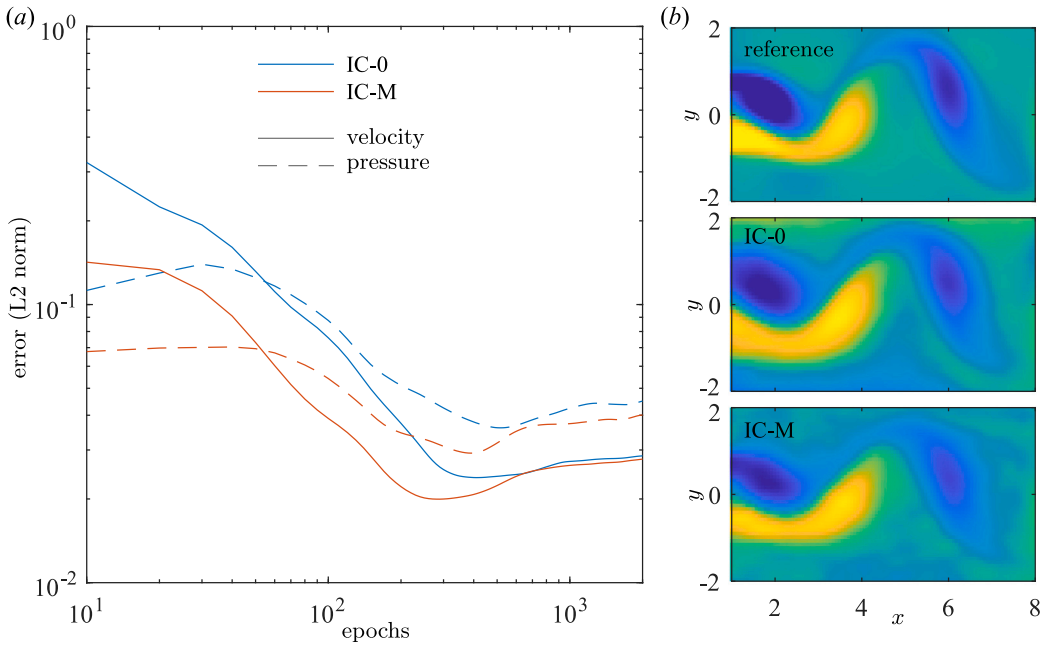


Fig. 9. Effect of initial condition on (*a*) convergence of the solution error, and (*b*) reconstructed vorticity fields. I0 uses all zeros for the initial velocity and pressure fields, while I1 uses the measurement data as an initial condition.

performance of both ODIL and PINN for reconstructing the velocity fields from low resolution and noisy 2D2C velocity measurements.

The test case considered in this section is the three-dimensional flow over a circular cylinder at $Re = 3900$. Numerical simulations were performed using an in-house spectral element solver, and the results are in agreement with Jiang and Cheng (2021). Fig. 17 plots isosurfaces of Q-criterion $Q = 1$ (normalised by the cylinder diameter and freestream velocity), showing the wake behind the cylinder. A transient animation is also provided in the online supplementary material. This flow features periodic shedding of large-scale wake vortices, which are primarily aligned in the spanwise direction. The flow is turbulent, featuring many small-scale three-dimensional structures. Fig. 18 plots the time-histories of the lift and drag coefficients, as well as their frequency spectra (power spectral densities). There is a clear dominant

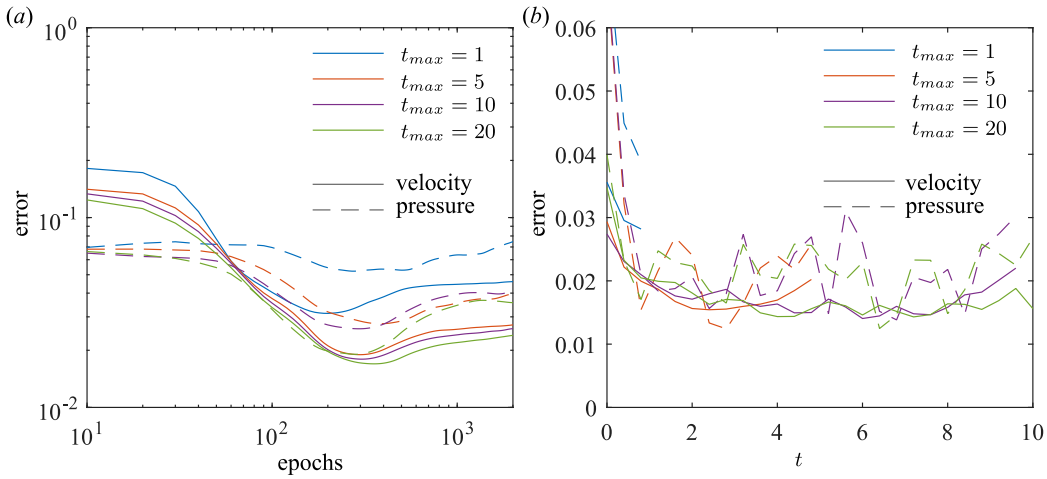


Fig. 10. (*a*) Effect of the maximum time (t_{max}) on the convergence history of ODIL, showing solution error against number of training epochs. (*b*) Variation of solution error against non-dimensional flow-time t, for different values of t_{max}.

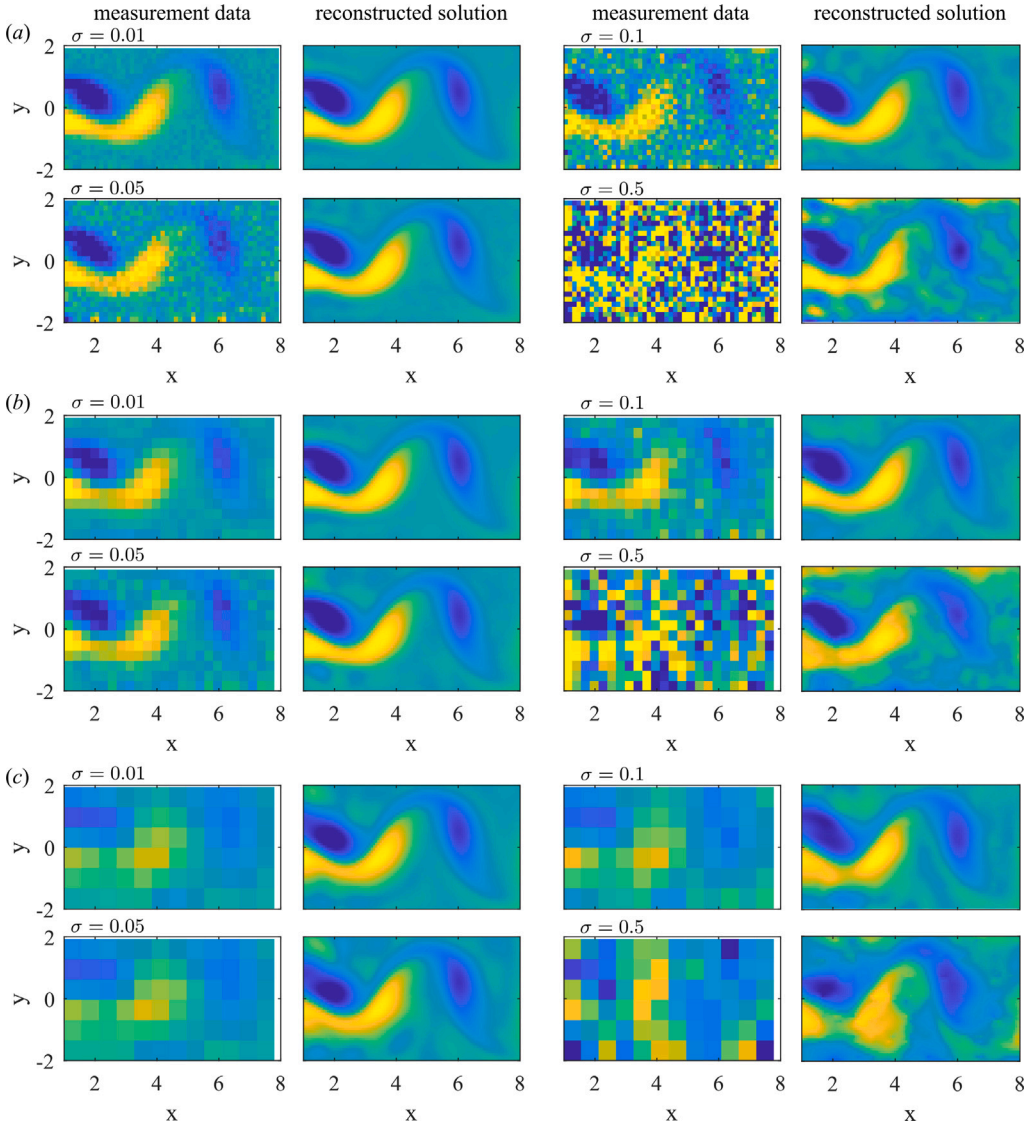


Fig. 11. Reconstructed vorticity fields for various noise levels (σ_ϵ) and data resolutions (*a*) $\Delta x = 0.283, \Delta y = 0.163$ and $\Delta t = 0.2$ (*b*) $\Delta x = 0.283, \Delta y = 0.327$ and $\Delta t = 0.4$, (*c*) $\Delta x = 0.566, \Delta y = 0.645$ and $\Delta t = 0.8$. An animated version of this figure is available in the online supplementary material. Each subfigure shows vorticity contours obtained from both the noisy measurement data (left) and the reconstructed solution (right).

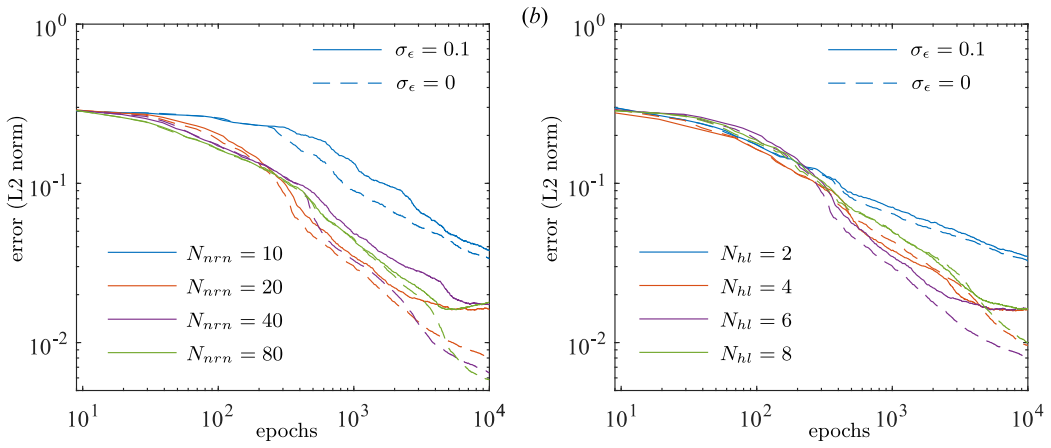


Fig. 12. PINN convergence history for different numbers of hidden layer (N_{hl}) and numbers of neurons per hidden layer (N_{nrn}). (*a*) Various N_{nrn} with $N_{hl} = 6$, and (*b*) various N_{hl} with $N_{nrn} = 20$.

Table 4
$L2$ norm of the error between the PINN reconstructions and the high-resolution reference solution for various resolution of the measurement data, and with added noise $\sigma_\epsilon = 0.1$. The number of epochs and training time required to achieve the same solution error as ODIL for a given measurement data are also provided.

Δx	Δy	Δt	σ_ϵ	E	E_u	E_p	Epochs	Walltime (s)
0.0707	0.0816	0.1	0.10	7.54×10^{-3}	8.13×10^{-3}	6.96×10^{-3}	–	–
0.1414	0.1633	0.2	0.10	8.94×10^{-3}	9.51×10^{-3}	8.37×10^{-3}	4836	733
0.2828	0.3265	0.4	0.10	1.41×10^{-2}	1.62×10^{-2}	1.20×10^{-2}	1180	161
0.5657	0.6531	0.8	0.10	3.27×10^{-2}	3.33×10^{-2}	3.20×10^{-2}	1766	267

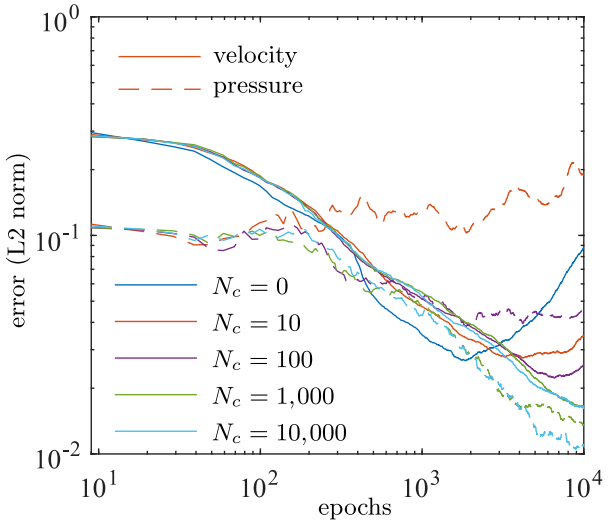


Fig. 13. Effects of the number of interior collocation points (N_c) on convergence history for PINN.

frequency with Strouhal number $St = 0.204$, which is the frequency of vortex shedding. However, broad-band noise is also present due to turbulence.

Synthetic measurement data are obtained by sampling the x and y velocity components from three-dimensional spectral-element solution, in a two-dimensional x–y plane. The $z = 0$ plane was selected for this study, although the precise z-location is not physically significant due to symmetry. Synthetic measurements with various spatial and temporal resolutions were produced, and Gaussian noise with standard deviation σ_ϵ was added. The domain of the measurement data was restricted to $x_{min} = 0.5$, $x_{max} = 3$, $y_{min} = -1$, $y_{max} = 1$, $t_{min} = 0$ and $t_{max} = 1$. The ODIL and PINN techniques were then used to reconstruct the high resolution x and y velocity components from the low-resolution and noisy measurements.

Since the pressure force, out-of-plane velocity component and out-of-plane derivatives are not available, these terms are modelled as an external force, so that the ith component of the momentum equation is written as:

$$\frac{\partial u_i}{\partial t} + u_x \frac{\partial u_i}{\partial x} + u_y \frac{\partial u_1}{\partial y} = f_i + \frac{1}{Re}\left(\frac{\partial^2 u_i}{\partial x^2} + \frac{\partial^2 u_i}{\partial y^2}\right), \tag{14}$$

$$f_i = -\frac{\partial p}{\partial x_i} - u_z \frac{\partial u_i}{\partial z} + \frac{1}{Re}\frac{\partial^2 u_i}{\partial z^2}. \tag{15}$$

The unknown force components f_x and f_y are included in the ODIL and PINN models, in addition to the velocity components u_x and u_y. To avoid over-fitting to added noise, the force components are regularised as follows:

$$L_{reg,f} = \left|\left|\nabla^2_{2D} \frac{\mathrm{d}f_x}{\mathrm{d}t}\right|\right|_2 + \left|\left|\nabla^2_{2D} \frac{\mathrm{d}f_y}{\mathrm{d}t}\right|\right|_2, \tag{16}$$

where $\mathrm{d}/\mathrm{d}t$ is the material derivative, and ∇_{2D} is the 2D gradient operator in the x-y plane. This regularisation term smooths the material derivative of the forcing term, promoting forcing terms that are approximately advected with the fluid. Additionally, it was found that the ODIL reconstruction resulted in unphysically large values of the term $\partial u_x/\partial x + \partial u_y/\partial y$ (the 2D velocity divergence). Therefore, additional regularisation was applied to this term:

$$L_{reg,u} = \left|\left|\nabla^2\left(\frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y}\right)\right|\right|_2. \tag{17}$$

For the 2D2C case considered here, physics information for the transport of momentum in the z-direction is not available, which will limit the accuracy of reconstruction compared to the fully two-dimensional case considered in the previous section. However, physics information for the transport of momentum in the x and y-directions (such as the terms $u_x \partial u_i/\partial x$) is available. In this section, we show that incorporating this limited physics information allows ODIL and PINN to produce more accurate flow reconstructions than conventional techniques such as Fourier filtering, which do not utilise this limited physics information. Moreover, only the x and y components of velocity

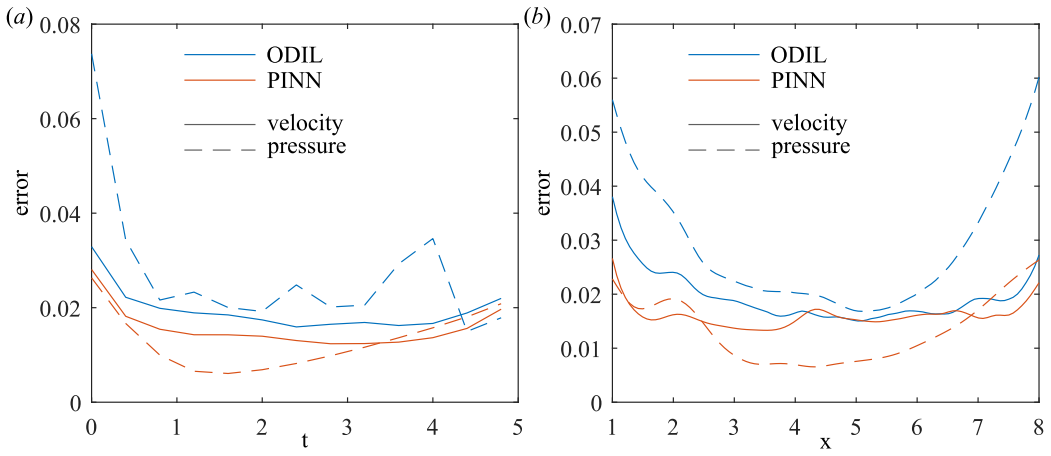


Fig. 14. A comparison between the distribution of solution error against (*a*) flow time t and (*b*) streamwise position x for both ODIL and PINN.

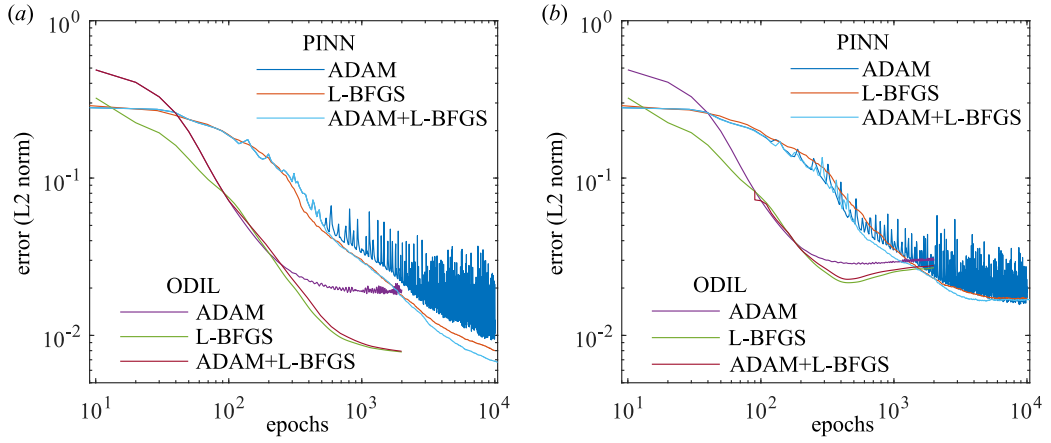


Fig. 15. Convergence histories for both ODIL and PINN using the ADAM and L-BFGS optimisers, for either (*a*) noiseless ($\sigma = 0$) or (*b*) noisy ($\sigma = 0.1$) data. The learning rate for ADAM was set to 1×10^{-2}.

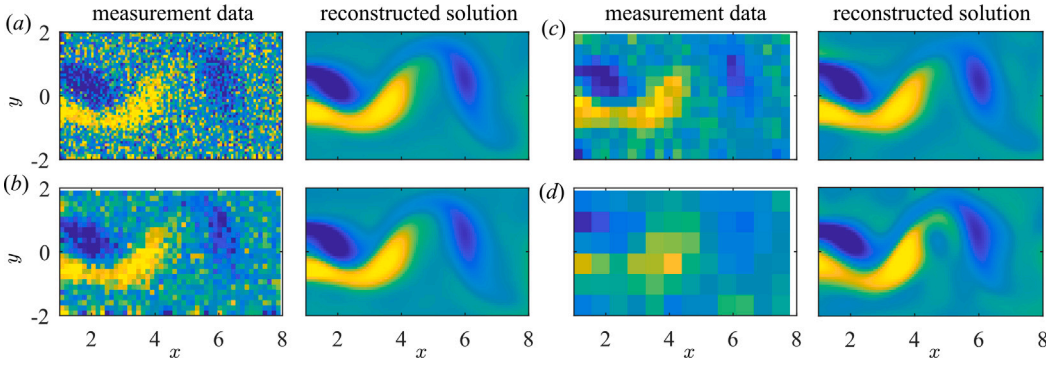


Fig. 16. Vorticity fields reconstructed using PINN for various data resolutions and with added noise $\sigma_\epsilon = 0.1$. The data resolutions are (*a*) $\Delta x = 0.0707$, $\Delta y = 0.0816$ and $\Delta t = 0.1$, (*b*) $\Delta x = 0.1414$, $\Delta y = 0.1633$ and $\Delta t = 0.2$, (*c*) $\Delta x = 0.2828$, $\Delta y = 0.3265$ and $\Delta t = 0.4$ and (*d*)$\Delta x = 0.5657$, $\Delta y = 0.6531$ and $\Delta t = 0.8$. Each subfigure shows vorticity contours obtained from both the noisy measurement data (left) and the reconstructed solution (right).

(and hence, the z-component of vorticity) may be recovered from 2D2C data. The out-of-plane (z) velocity, and hence the in-plane (x and y vorticity) vorticity components are unable to be reconstructed.

Synthetic datasets were generated with a range of spatial resolutions and added noise (σ_ϵ). Fig. 19(*b*) shows the high-resolution reference solution, while subfigure (*a*) shows an example synthetic measurement data, with $\Delta x = \Delta y = 0.04$ and $\sigma_\epsilon = 0.1$. Reconstructed vorticity fields using both PINN and ODIL are presented in subfigures (*c*) and (*d*), respectively. The settings for ODIL were selected to minimise solution error, as discussed in Section 4.1. For PINN, a neural network with 6 hidden layers and 80 neurons per hidden layer was used, and the number of interior collocation points was 10,000. These settings were the largest neural network that could be used given the memory limitations of the NVIDIA GeForce RTX 4090 graphics card. Additionally, $W_{phys} = 1$ and $W_{dat} = 1$ were used for PINN, which achieves fastest convergence, and no regularisation terms were needed. Table 5 lists the hyperparameters used for both ODIL and PINN for this case.

Fig. 19(*c*) shows that PINN successfully removes noise, and can resolve large and medium-scale flow structures. However, the solution is overly smoothed and PINN does not reconstruct small-scale flow

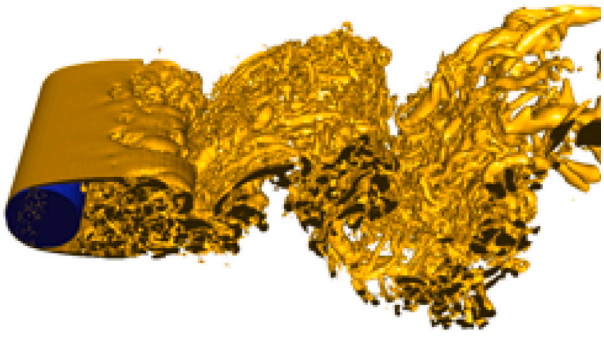


Fig. 17. Visualisation of the wake behind a circular cylinder at $Re = 3900$, showing isosurfaces of Q-criterion $Q = 1$. Flow is from left to right.

Table 5
List of hyperparameters used for both ODIL and PINN to reconstruct synthetic 2D2C measurements of the 3D cylinder wake.

ODIL		PINN	
W_{meas}	1	W_{meas}	1
W_{phys}	2	W_{phys}	1
$W_{\nabla\cdot u}$	10^{-4}	N_{hl}	6
W_f	10^{-6}	N_{nrn}	80
Δx_r	0.01	N_c	10,000
Δy_r	0.01	l_r	10^{-2}
Δt_r	0.0125	N_{adam}	5000
m_{lbfgs}	5	m_{lbfgs}	5

structures. Subfigure (*d*) shows that ODIL resolves some smaller-scale structures. However, ODIL also contains significant measurement noise, which appears to have a similar length-scale to the small-scale flow vortical structures in the reference solution.

Fig. 20 compares the convergence histories of ODIL and PINN for the case with data resolution $\Delta x = \Delta y = 0.04$ and $\sigma_\epsilon = 0.1$. For this comparison, ODIL was trained using the L-BFGS optimiser, while PINN was trained using the ADAM optimiser for the first 5000 epochs before switching to the L-BFGS optimiser. ODIL initially converges faster than PINN, and reaches a minimum error of 5.29×10^{-2} in 3650 epochs, requiring 1358 s (22 min) of walltime. After further training, the error increases due to over-fitting to measurement noise. PINN converges much slower than ODIL, requiring 30,320 epochs of training (40,200 s or 11.2 h of walltime) to reduce the solution error to below 5.29×10^{-2}. However, PINN evidently does not suffer from overfitting, and the solution error continues to decrease with further training time. After 60,000 epochs of training (requiring 81,800 s or 22.7 h of walltime), the error is reduced to 4.67×10^{-2}. Fig. 20 suggests that the PINN error would continue to decrease with further training time. However it would take multiple days of training time to achieve significant further improvements.

Hence, although ODIL converges much faster than PINN, ODIL is more prone to overfitting, and a lower overall solution error may be achieved using PINN. However, PINN suffers from long training times, which may make it impractical for large scale problems. Since ODIL and PINN rely on the same physics-informed principles, this indicates that the success of PINN is not only due to the use of physics. The fact that PINNs do not overfit to measurement data is likely a result of spectral bias (Wang et al., 2021), where neural networks more easily learn the low-frequency (long-wavelength) structures (i.e. the large vortices) than high frequency structures (small vortices and added noise). The spectral bias of PINNs also explains why regularisation terms were not required to smooth the solutions, as was needed for ODIL. Various techniques, such as Fourier feature embeddings (Tancik et al., 2020), have been proposed to reduce spectral bias when training neural networks. This may improve the training time of PINNs, by fitting to all frequency components equally. However, the results of the preset paper suggest that this might result in overfitting to measurement noise, rather than an improved reconstruction. This requires further investigation.

4.1. ODIL

We now discuss the effects of various loss-function weights in the performance of ODIL. These weights include the physics weight W_{phys}, data weight W_{dat}, and the regularisation weights for the forcing term W_f and the velocity divergence term $W_{\nabla\cdot u}$. The optimal weights were found to be $W_{phys} = 2$, $W_{dat} = 1$, $W_f = 10^{-6}$ and $W_{\nabla\cdot u} = 10^{-4}$. Unless otherwise stated, these parameters are used in the remainder of this section. The reconstruction resolution was kept constant ($\Delta x = 0.01$, $\Delta y = 0.01$ and $\Delta t = 0.0125$), with the spatial resolution chosen to match the high-resolution measurement data, and the temporal resolution was chosen to keep the maximum Courant number close to unity. Finally, the L2 norm was used for all terms in the loss function and the measurement data was used as an initial condition, as these settings were found to achieve best performance in the previous case (Section 3).

Fig. 21 plots the effect of physics weight W_{phys} on the (*a*) convergence history of solution error and (*b*) vorticity plots of the reconstructed solution. The maximum solution error is relatively insensitive to the physics weight, and the reconstructed vorticity fields are similar regardless of physics weight. However, the training time is reduced for reduced physics weight. The lowest solution error is achieved for $W_{phys} = 2$, and therefore $W_{phys} = 2$ is used for the remainder of this section.

Fig. 22(*a*) plots the effect of the weight for the velocity divergence regularisation term ($W_{\nabla\cdot u}$) on the convergence history of the solution error. The lowest solution error is achieved for $W_{\nabla\cdot u} = 10^{-4}$, and this is used for the remainder of this section. Fig. 23(*a*) plots both the vorticity field and the 2D velocity divergence ($\nabla_{2D}\cdot\boldsymbol{u} = \partial u/\partial x + \partial v/\partial y$) obtained for various values of $W_{\nabla\cdot u}$. For no regularisation ($W_{\nabla\cdot u} = 0$), unphysically large velocity divergences are observed, resulting in thinner vortex structures. The regularisation term prevents large values of 2D velocity divergence from forming, resulting in smoother vortex structures in the reconstructed field, and achieving a lower solution error.

Fig. 22(*b*) plots the effect of the weight for the force regularisation term (W_f) on the convergence history of the solution error. A small regularisation ($W_f = 10^{-6}$) produces a small improvement in the solution compared to the solution with no regularisation. Further increasing the regularisation to $W_f = 10^{-5}$ results in a slightly less accurate solution. However, increasing W_f reduces the impact of overfitting on the solution error.

Fig. 24 plots vorticity fields and the x-component of the out-of-plane forcing terms (f_x), obtained using various force regularisation weights W_f. Subfigure (*a*) corresponds to the number of training epochs for which the solution error is minimised, while subfigure (*b*) corresponds to 20,000 training epochs. Increasing W_f results in a smoother out-of-plane forcing field, resulting in less over-fitting to noise. This is particularly notable after a large number of training epochs, where solutions with a low W_f substantially overfit to the measurement noise, while $W_f = 10^{-5}$ produces a much more accurate solution. However, $W_f = 10^{-6}$ produces a lower minimum solution error, and is used for the remainder of this study.

Fig. 25 plots the vorticity fields reconstructed using ODIL for various noise levels and data resolutions. The corresponding solution errors, number of training epochs and walltime used are provided in Table 6. The accuracy of the ODIL reconstruction increases as either the level of noise in the measurement data is reduced, or the resolution of the data increases. Reasonable reconstructions are achieved for added noise up to $\sigma_\epsilon = 0.1$. For $\sigma_\epsilon = 0.5$, substantial added noise remains in the reconstructed vorticity fields, and the vortical structures in the wake are not clearly visible.

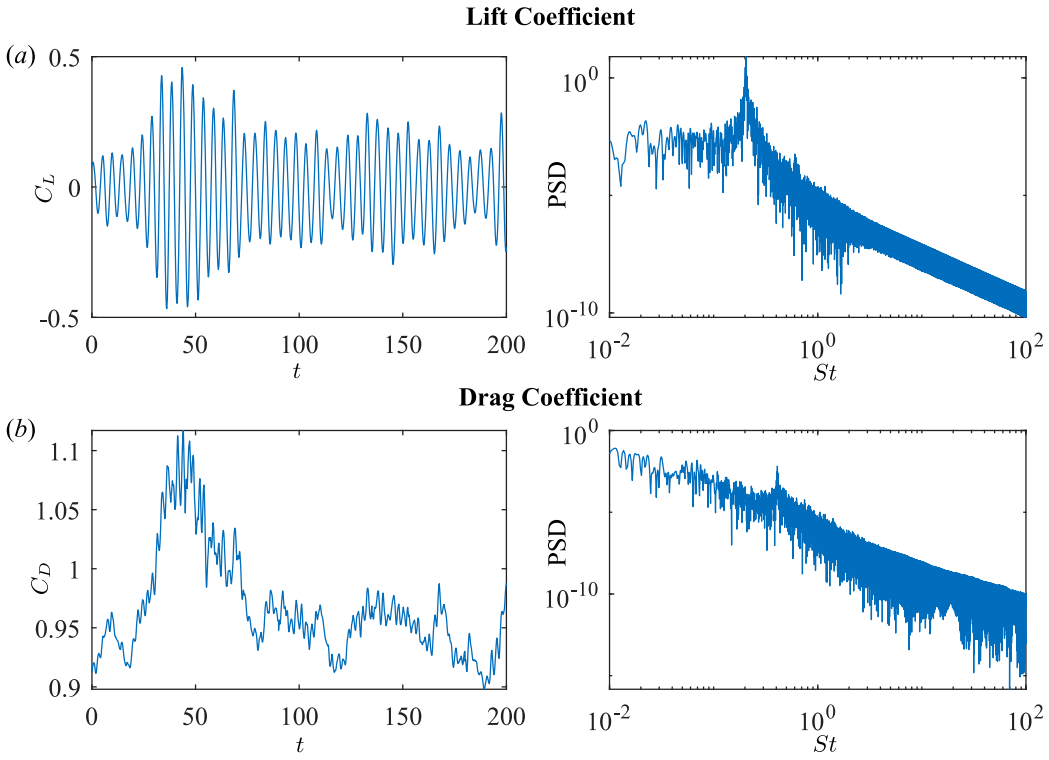


Fig. 18. Time histories (left) and power spectral densities (right) of the (*a*) lift coefficient and (*b*) drag coefficient.

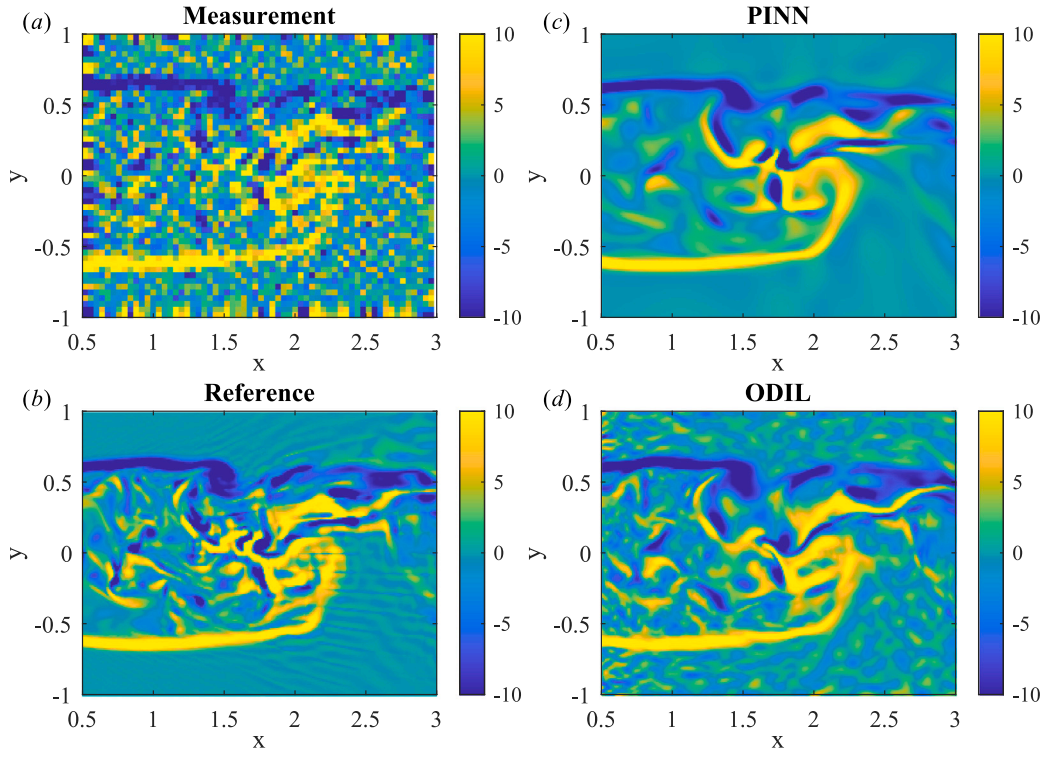


Fig. 19. Plots of spanwise (z) vorticity for (*a*) the synthetic measurement data, (*b*) the reference solution, (*c*) the PINN reconstruction and (*d*) the ODIL reconstruction, with measurement resolution $\Delta x = \Delta y = 0.08$ and $\Delta t = 0.1$. An animated version of this figure is available in the online supplementary material.

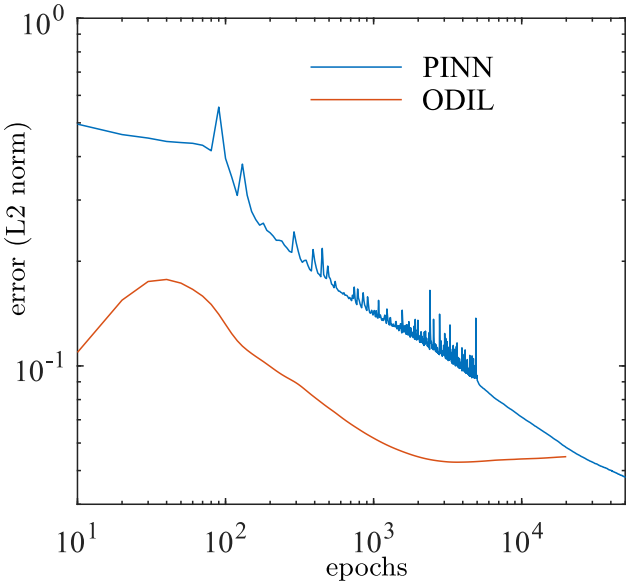


Fig. 20. Comparison between ODIL and PINN convergence histories, showing the L2 norm of error between the reconstructed velocity fields and the reference solution.

Table 6
L2 norm of the error between the ODIL reconstructions and the high-resolution reference solution, as well as the number of training epochs and training wall-time, for various resolution and noise of the measurement data.

Δx	Δy	Δt	σ_ϵ	Error (L2)	Epochs	Walltime (s)
0.0200	0.0200	0.1	0.01	2.82×10^{-2}	20 000	7488
			0.05	3.21×10^{-2}	20 000	7484
			0.10	3.90×10^{-2}	8850	3315
			0.50	1.16×10^{-1}	12 270	4598
0.0400	0.0400	0.1	0.01	3.23×10^{-2}	18 540	6872
			0.05	4.09×10^{-2}	6980	2585
			0.10	5.13×10^{-2}	3450	1283
			0.50	1.16×10^{-1}	1400	528
0.0800	0.0800	0.1	0.01	4.93×10^{-2}	5500	2235
			0.05	5.72×10^{-2}	3500	1437
			0.10	7.09×10^{-2}	1640	676
			0.50	1.41×10^{-1}	330	146

4.2. Comparison to low-pass filtering

In this subsection, the ODIL and PINN reconstructions are compared to the results of low-pass filtering, which represents a conventional statistical approach to noise removal. The synthetic measurement data (with $\Delta x = \Delta y = 0.04$, $\Delta t = 0.1$ and $\sigma_\epsilon = 0.1$) were transformed into the frequency domain by the fast-Fourier transform (FFT). A Gaussian low-pass filter was then applied to remove high-frequency noise:

$$G(\omega) = \exp(-(k_x^2 + k_y^2)/K_{filt}^2), \tag{18}$$

where k_x and k_y are the wavenumbers in the x and y directions, and K_{filt} is the low-pass frequency of the filter. A value of $K_{filt} = 0.13$ was found to give the lowest error. Reconstructed velocity fields were then obtained using the inverse FFT algorithm. During this process, the velocity fields were upscaled by padding with zeros in the high-frequencies.

Fig. 26 plots the vorticity field obtained using the Gaussian low-pass filter. The solution error obtained using the Gaussian filter was 8.17×10^{-2}, which is larger than the error obtained using ODIL (5.13×10^{-2}) and PINN (4.67×10^{-2}). Therefore, physics-informed approaches are more effective at de-noising and super-resolution than purely statistical approaches.

Fig. 27 presents a comparison between the ODIL, PINN and Fourier-filtered reconstructions in the frequency domain. Here, the power-spectral density of x-velocity was obtained by applying FFT in the x-direction only, then averaging over y and t. Similar results were obtained when using either the y-velocity component, or when applying FFT in the y-direction. The original measurement data contains significant high-frequency components, which represent the added noise. The measurement data does not contain frequency components with wavelengths less than the measurement resolution ($\Delta x = 0.04$). Fourier filtering removes the high-frequency noise. However, it also removes high-frequency components of the true solution. ODIL and PINN produce a similar frequency spectrum to each other. High frequency noise is mostly removed, however high frequencies corresponding to the reference solution remain. Moreover, high frequency components not originally present in the measurement data are reconstructed, indicating successful upscaling. However, both ODIL and PINN underestimate the energy contained in medium-wavelength structures with wavenumbers between 2 and 10 (wavelengths between 0.5 and 0.1). These wavelengths are comparable to small and medium-scale vortical structures in the wake, suggesting that ODIL and PINN slightly smooth out these structures when reconstructing the solution.

5. Experimental PIV data

Finally, both ODIL and PINN are applied to the reconstruction of experimental planar PIV velocity measurements. The PIV measurements were obtained in an experimental study into the flow-induced vibrations of a thin elliptical cylinder (Lo et al., 2023), performed in the free-surface recirculating water channel of the Fluids Laboratory for Aeronautical and Industrial Research (FLAIR) at Monash University.

The flow configuration is as shown in Fig. 28. A thin elliptical cylinder is mounted on a low-friction air-bearing rig and immersed in the water channel. This allows the cylinder to undergo single-degree-of-freedom flow-induced vibration in the direction perpendicular to the oncoming flow. The system behaves as a spring–mass–damper system, with the spring constant provided by a pair of precision extension springs and the damping coefficient provided by an adjustable eddy-current magnetic damper mechanism. For further details of the experimental procedure, refer to Lo et al. (2023) and the references therein.

The behaviour of this system is controlled by several parameters, including the aspect ratio $AR = b/a$, mass ratio $m^* = m/m_d$, damping constant $\xi = c/(2\sqrt{k(m+m_a)}$, Reynolds number $Re = Ub/\nu$ and reduced velocity $U^* = U/(f_{nw}b)$. Here, a and b are the length of the ellipse minor and major axes, respectively, m is the total oscillating mass, m_a is the added mass, m_d is the mass of fluid displaced by the cylinder, c is the structural damping coefficient, k is the spring coefficient, U is the freestream velocity and f_{nw} is the natural frequency of the system based on a pluck test in still water. The test case considered in the present study uses parameters $AR = 5$, $m^* = 17.4$, $U^* = 8.0$, $\xi =$ and $Re = 2746$. For these parameters, the cylinder oscillates with an amplitude of $y_{max}/D = 0.604$, and sheds a single primary vortex each half-cycle of the motion approximately at the point of maximum displacement.

A particle-induced velocimetry system was used to obtain a set of velocity-field snapshots for this flow, as described in Lo et al. (2023). The non-dimensional resolution of the measurement data is $\Delta x = \Delta y = 0.0933$, and $\Delta t = 0.439$, normalised by the cylinder major axis and the freestream velocity respectively. The measurement data snapshots were then cropped to reduce the domain size, with $x_{min} = 1$, $x_{max} = 5$, $y_{min} = -2.5$ and $y_{max} = 2.5$. A total of 12 snapshots were included in the measurement data, so that $t_{min} = 0$ and $t_{max} = 4.83$.

Fig. 29 presents a comparison between the vorticity computed directly from the PIV velocity fields, as well as both the ODIL and PINN reconstructions. The results are qualitatively similar to the comparison presented previously in Fig. 19. Namely, the ODIL reconstruction still contains measurement noise, while the PINN reconstruction likely provides excessive smoothing, and does not reconstruct small-scale structures. The hyperparameters used for this case are listed in Table 7.

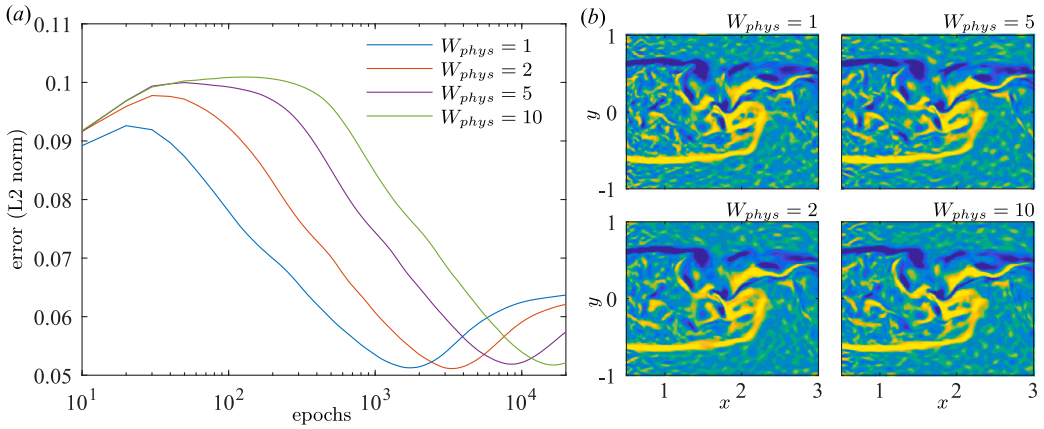


Fig. 21. Effect of physics weight on (*a*) convergence of the solution error, and (*b*) reconstructed vorticity fields.

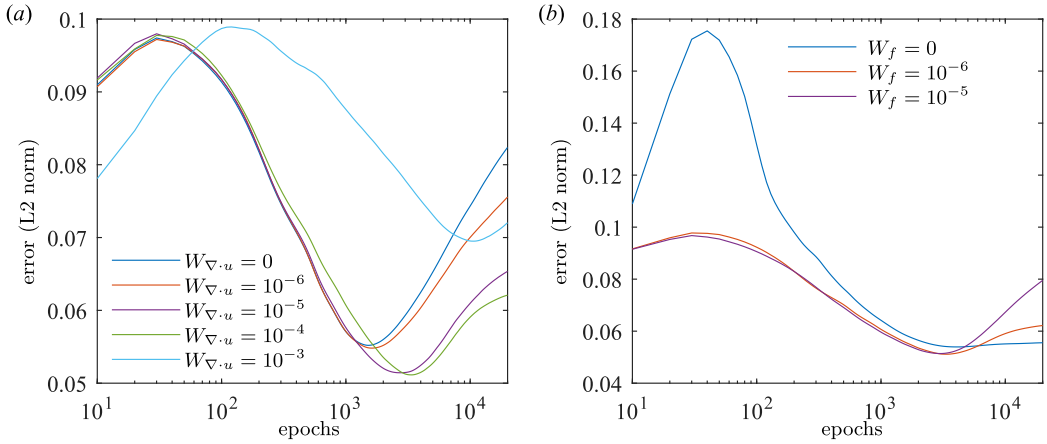


Fig. 22. Effect of regularisation weights (*a*) $W_{\nabla\cdot u}$ and (*b*) W_f on the convergence history of the L2 norm of error between the reconstructed velocity fields and the reference solution.

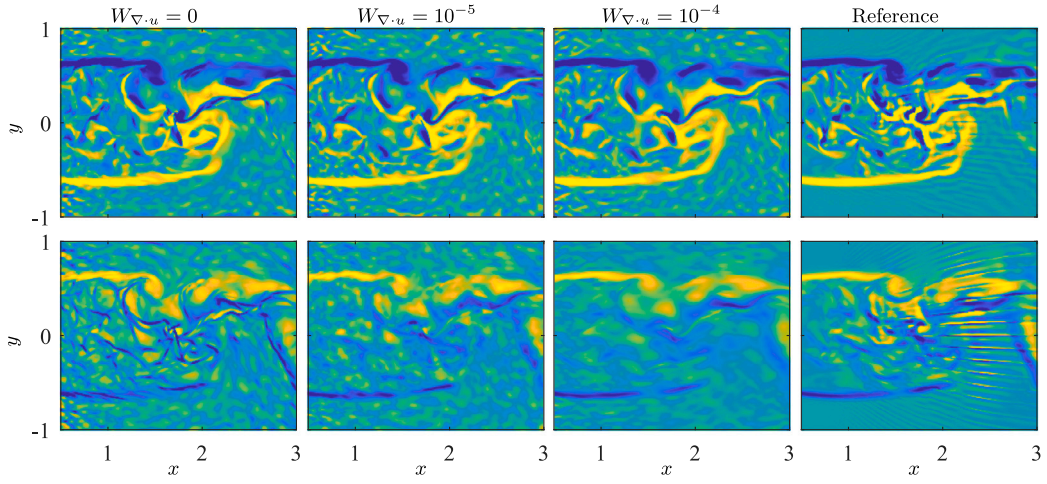


Fig. 23. Effect of the velocity regularisation weight $W_{\nabla\cdot u}$ on the vorticity field (top) and 2D velocity divergence (bottom).

Fig. 30(*a–c*) plots the convergence history using ODIL for various values of the physics and regularisation weights. Since the true solution is not known, we plot the L2 norm of error between the validation data (10% of the measurement data) and the reconstructed solution. The general trends are comparable to those seen in Figs. 21 and 22. The optimal weights producing minimum validation error are found to be $W_{\nabla\cdot u} = 10^{-4}$, $W_f = 10^{-5}$ and $W_{phys} = 1$.

Fig. 30(*d*) plots the convergence histories for both ODIL and PINN. Once again, ODIL was trained using the L-BFGS optimiser for all epochs, while PINN was trained using the ADAM optimiser for the first 5000 epochs before switching to the L-BFGS optimiser. As seen in previous sections, ODIL converges much faster than PINN, requiring 1520 training epochs (482 s of walltime) to reach a minimum validation error of 5.49×10^{-2}. PINN converges much slower than ODIL, reaching a

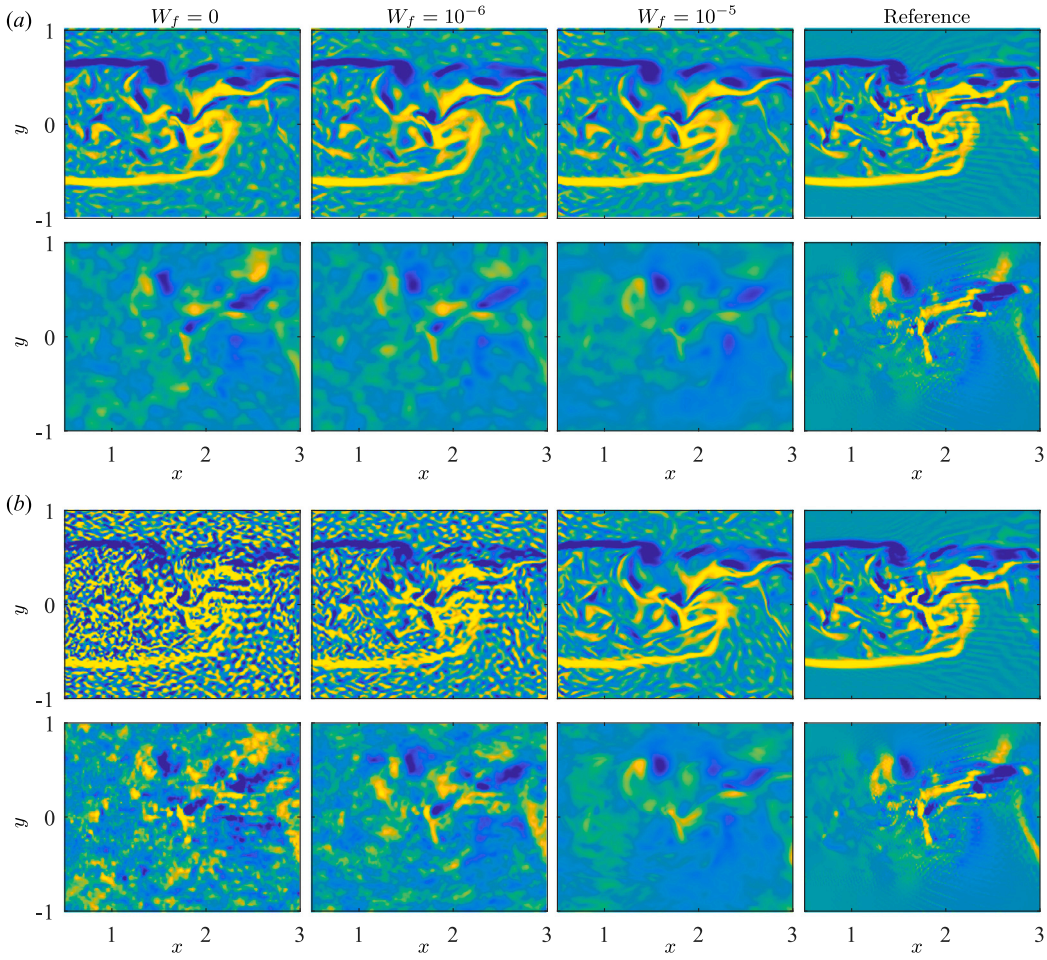


Fig. 24. Effect of the force regularisation weight W_f on the vorticity field (top) and out-of-plane forcing term f_x (bottom). (*a*) corresponds to optimally trained solution (number of epochs corresponding to lowest solution error), while (*b*) corresponds to the over-trained solution (20,000 training epochs).

Table 7
List of hyperparameters used for both ODIL and PINN to reconstruct experimental PIV data.

ODIL		PINN	
W_{meas}	1	W_{meas}	1
W_{phys}	1	W_{phys}	1
$W_{\nabla \cdot u}$	10^{-4}	N_{hl}	6
W_f	10^{-5}	N_{nrn}	80
Δx_r	0.0233	N_c	10,000
Δy_r	0.0233	l_r	10^{-2}
Δt_r	0.0549	N_{adam}	5000
m_{lbfgs}	5	m_{lbfgs}	5

minimum error of 5.60×10^{-2} after 39,600 epochs, requiring 50,600 s (14.1 h) of walltime.

6. Conclusions

We compared two techniques – optimising a discrete loss (ODIL), and physics informed neural networks (PINN) – for the reconstruction of velocity fields from low resolution and noisy planar PIV measurements. Both techniques can accurately reconstruct the velocity and pressure from low resolution and noisy measurements sampled from a two-dimensional flow field. When data are provided for only two velocity components in a two-dimensional plane, sampled from a three-dimensional flow (i.e. planar PIV measurements), ODIL can produce reasonable reconstructions of the velocity field when the measurement noise is low. However, ODIL reconstructions do not eliminate measurement noise from the data. The PINN reconstruction is less prone to overfitting, and can achieve a more accurate reconstruction. However, PINN also may not reconstruct small-scale flow structures. In all cases, ODIL requires fewer epochs of training time and less wall-time than PINN to produce results of similar accuracy.

CRediT authorship contribution statement

Stephen Terrington: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Conceptualization, Data curation, Formal analysis. **Mark Thompson:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization, Data curation, Software, Visualization, Project administration, Resources. **Kerry Hourigan:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization, Project administration, Resources.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (projects DP210100990 and DP190103388).

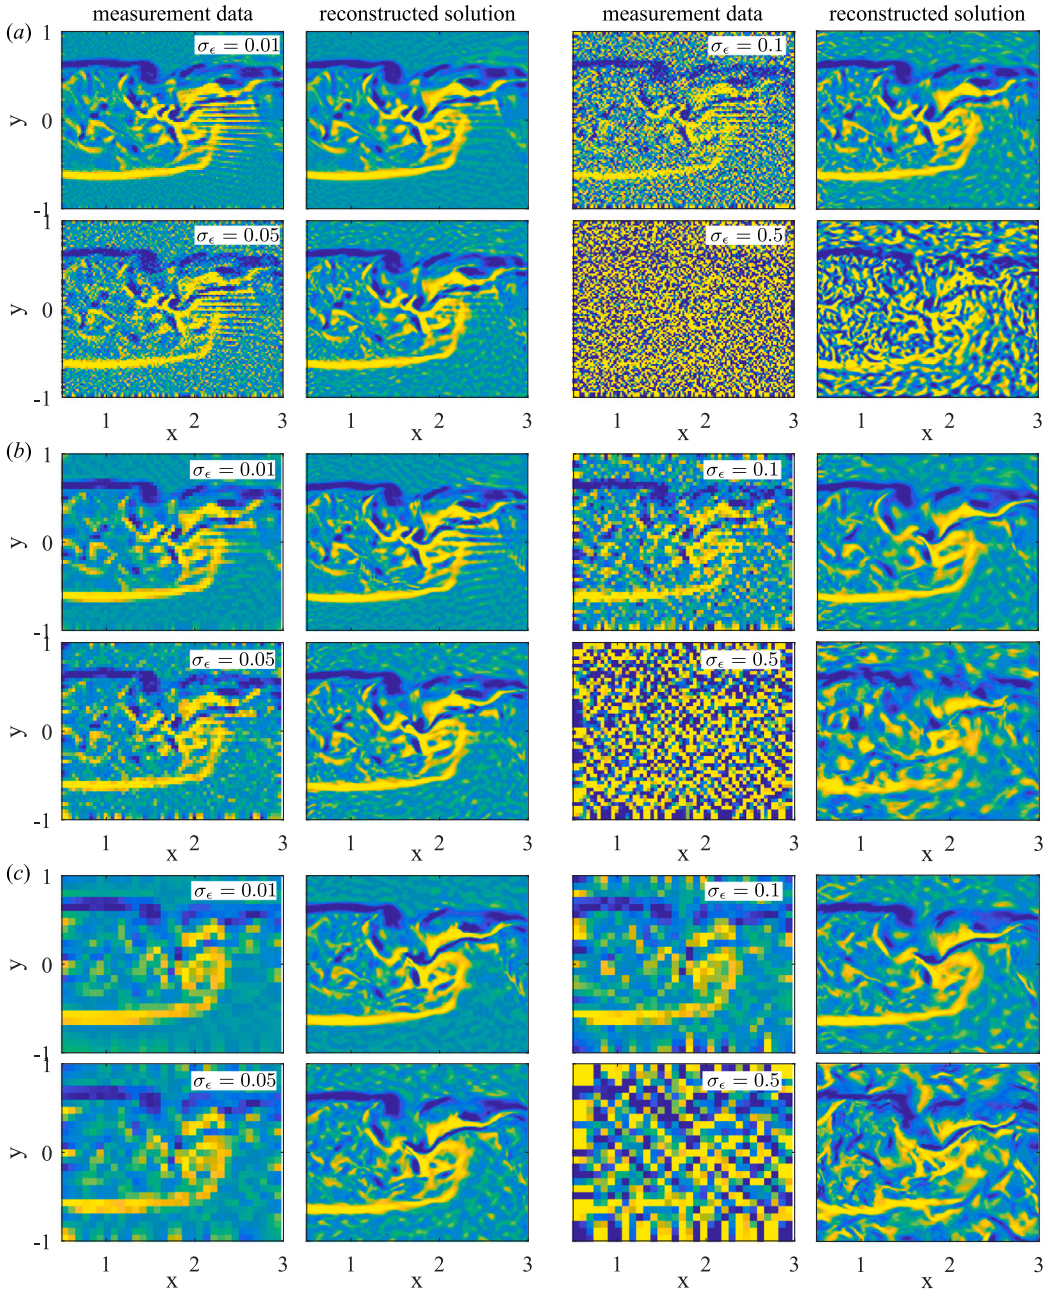


Fig. 25. Vorticity fields reconstructed using ODIL for various noise levels (σ_ϵ) and data resolutions (*a*) $\Delta x = \Delta y = 0.02$ and $\Delta t = 0.1$, (*b*) $\Delta x = \Delta y = 0.04$ and $\Delta t = 0.1$ and (*c*) $\Delta x = \Delta y = 0.08$ and $\Delta t = 0.1$. An animated version of this figure is available in the online supplementary material. Each subfigure shows vorticity contours obtained from both the noisy measurement data (left) and the reconstructed solution (right).

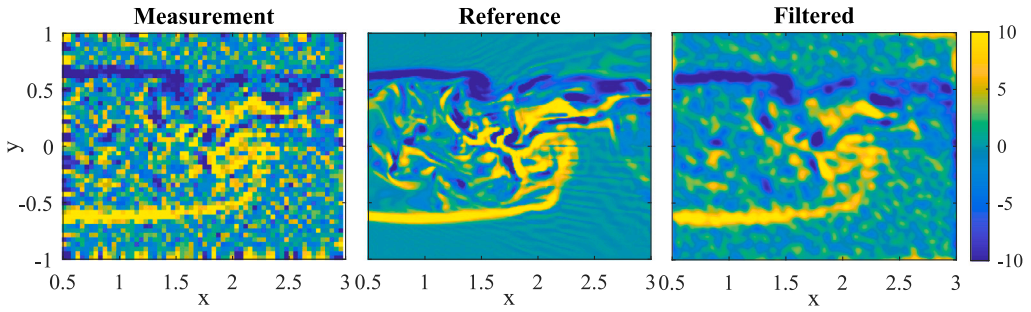


Fig. 26. Vorticity field obtained by applying a Gaussian low-pass filter to the velocity fields in the measurement plane.

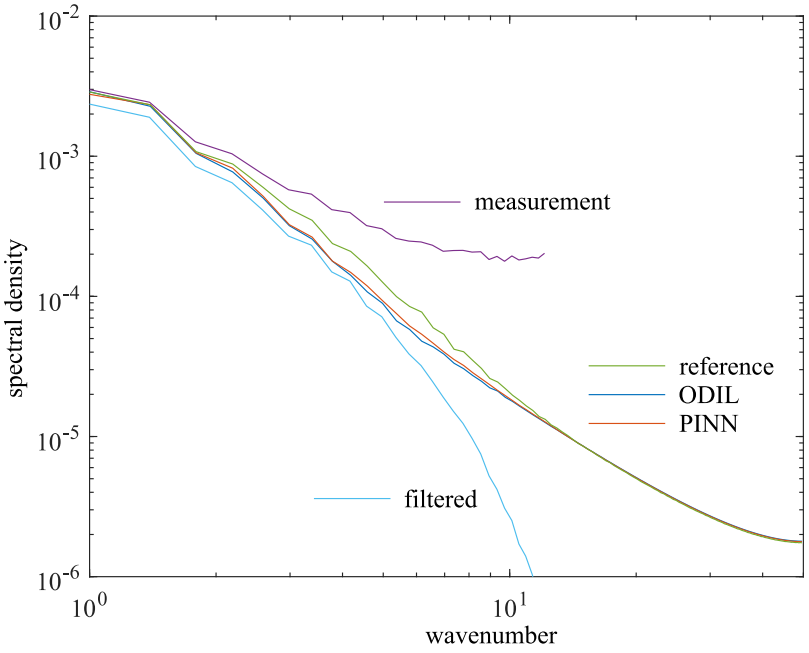


Fig. 27. Comparison between the power spectral densities of the measurement data, the reference solution, and the reconstructions obtained using ODIL, PINN and Gaussian low-pass filtering.

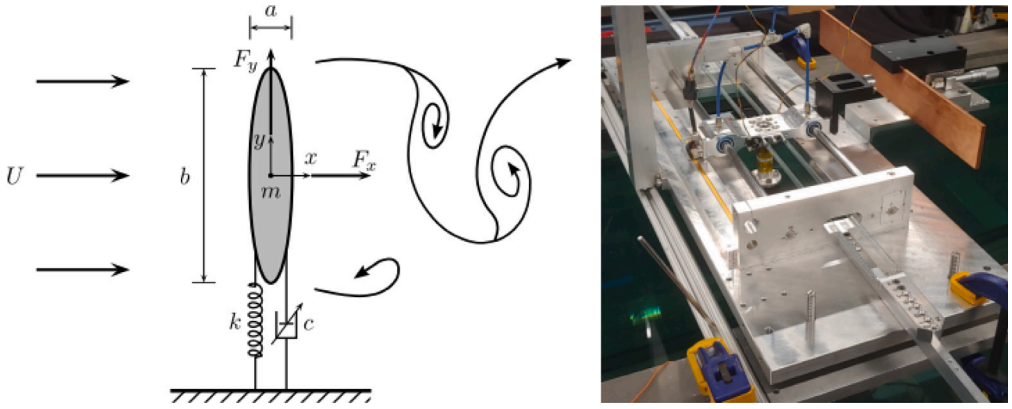


Fig. 28. Schematic illustration (left) and photograph (right) of the experimental configuration. An elliptical cylinder is mounted in a water channel using a low-friction air-bearing rig, which acts as a spring–mass–damper system.
Source: Reproduced with permission from Lo et al. (2023).

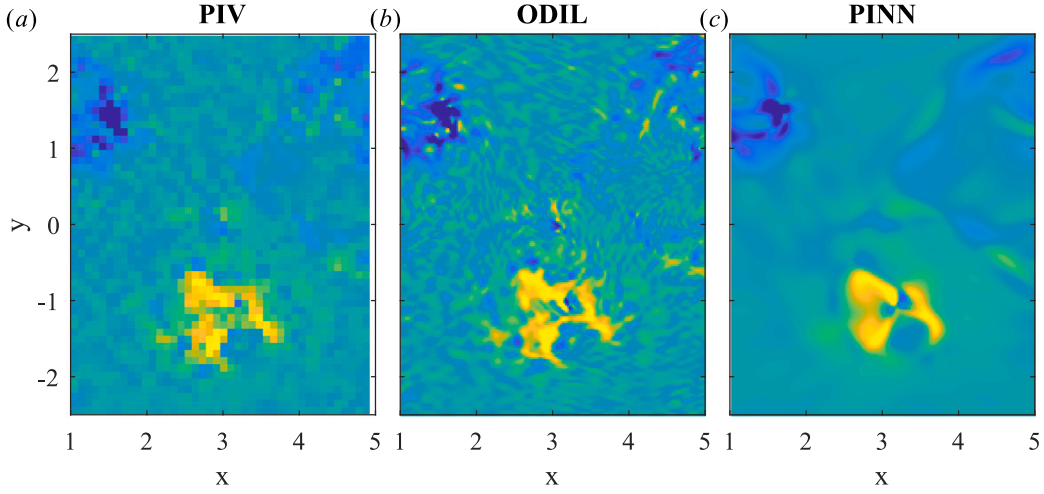


Fig. 29. Comparison between the vorticity fields for (*a*) the PIV measurement, (*b*) the ODIL reconstruction and (*c*) the PINN reconstruction. An animated version of this figure is available in the online supplementary material.

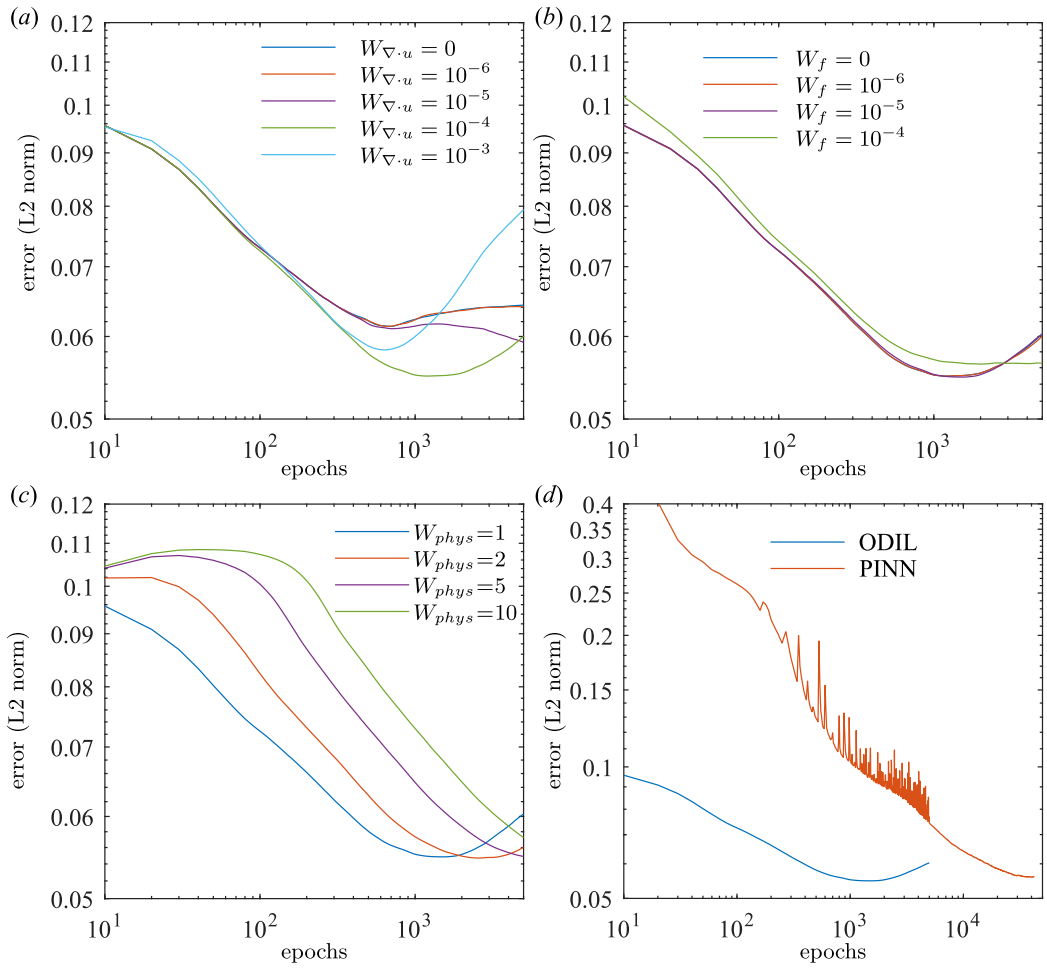


Fig. 30. Effect of (*a*) velocity regularisation weight $W_{\nabla\cdot u}$, (*b*) force regularisation weight W_f and (*c*) physics weight W_{phys} on the convergence history using ODIL. (*d*) Comparison of ODIL and PINN convergence histories. The error provided here is the L2 norm of error between the validation data and the reconstructed solution.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ijheatfluidflow.2025.109988.

Data availability

Data will be made available on request.

References

Fan, D., Xu, Y., Wang, H., Wang, J., 2023. Comparative assessment for pressure field reconstruction based on physics-informed neural network. Phys. Fluids 35, 077116.

Fukami, K., Fukagata, K., Taira, K., 2023. Super-resolution analysis via machine learning: a survey for fluid flows. Theor. Comput. Fluid Dyn. 37 (4), 421–444.

Gillissen, J.J., Bouffanais, R., Yue, D.K., 2019. Data assimilation method to de-noise and de-filter particle image velocimetry data. J. Fluid Mech. 877, 196–213.

Hasanuzzaman, G., Eivazi, H., Merbold, S., Egbers, C., Vinuesa, R., 2023. Enhancement of PIV measurements via physics-informed neural networks. Meas. Sci. Technol. 34, 044002.

He, C., Liu, Y., Gan, L., 2018. A data assimilation model for turbulent flows using continuous adjoint formulation. Phys. Fluids 30 (10), 105108.

He, C., Wang, P., Liu, Y., Gan, L., 2022. Flow enhancement of tomographic particle image velocimetry measurements using sequential data assimilation. Phys. Fluids 34 (3), 035101.

Jiang, H., Cheng, L., 2021. Large-eddy simulation of flow past a circular cylinder for Reynolds numbers 400 to 3900. Phys. Fluids 33, 034119.

Karnakov, P., Litvinov, S., Koumoutsakos, P., 2023. Flow reconstruction by multiresolution optimization of a discrete loss with automatic differentiation. Eur. Phys. J. E 46, 59.

Lo, J.C.C., Hourigan, K., Thompson, M., Zhao, J., 2023. The effect of structural damping on flow-induced vibration of a thin elliptical cylinder. J. Fluid Mech. 974, A5.

Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J. Comput. Phys. 378, 686–707.

Sofos, F., Drikakis, D., 2025. A review of deep learning for super-resolution in fluid flows. Phys. Fluids 37 (4), 041303.

Sofos, F., Drikakis, D., Kokkinakis, I.W., 2024. Comparison of super-resolution deep learning models for flow imaging. Comp. Fluids 283, 106396.

Steinfurth, B., Wiess, J., 2024. Assimilating experimental data of a mean three-dimensional separated flow using physics-informed neural networks. Phys. Fluids 36, 015131.

Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R., 2020. Fourier features let networks learn high frequency functions in low dimensional domains. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (Eds.), NIPS'20: Proceedings of the 34th International Conference on Neural Information Processing Systems. Curran Associates Inc, pp. 7537–7547.

Vétel, J., Garon, A., Pelletier, D., 2011. Denoising methods for time-resolved PIV measurements. Exp. Fluids 51, 893–916.

Wang, H., Liu, Y., Wang, S., 2022. Dense velocity reconstruction from particle image velocimetry/particle tracking velocimetry using a physics-informed neural network. Phys. Fluids 34, 017116.

Wang, S., Wang, H., Perdikaris, P., 2021. On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks. Comp. Methods Appl. Mech. Eng. 384, 113938.

Xu, S., Sun, Z., Huang, R., Guo, D., Yang, G., Ju, S., 2022. A practical approach to flow field reconstruction with sparse or incomplete data through physics informed neural network. Acta Mech. Sin. 39, 322302.